Playing Games with Bounded Entropy

Mehrdad Valizadeh, Amin Gohari

Department of Electrical Engineering, Sharif University of Technology, Tehran, Iran
(valizadeh@ee.sharif.edu, aminzadeh@sharif.edu)

arXiv:1702.05719v2 [cs.IT] 17 Apr 2017

Abstract

We study a two-player zero-sum game in which one of the players is restricted to mixed strategies with limited randomness. More precisely, we consider the maximum payoff that the maximizer (Alice) can secure with limited randomness h. This problem finds an operational interpretation in the context of repeated games with non-ideal sources of randomness as shown by Gossner and Vieille, and Neyman and Okada. We begin by simplifying the proof of Gossner and Vieille and also generalize their result. Then, we turn to the computational aspect of the problem, which has not received much attention in the game theory literature. We observe the equivalence of this problem with entropy minimization problems in other scientific contexts. Next, we provide two explicit lower bounds on the entropy-payoff tradeoff curve. To do this, we provide and utilize new results for the set of distribution that guarantee a certain payoff for Alice (mixed strategies corresponding to a security level for Alice). In particular, we study how this set of distribution shrinks as we increase the security level. While the use of total variation distance is common in game theory, our derivation indicates the suitability of utilizing the Renyi-divergence of order two.[1]

Keywords:

Repeated Games, Bounded Entropy, Entropy Minimization, Information Theory

[1]A short version of this paper has been accepted for presentation at the 2017 IEEE Symposium On Information Theory (ISIT 2017).

1. Introduction

Consider a two-player and zero-sum game between Alice and Bob. The game is played once. A Nash equilibrium assigns a strategy to Alice and Bob such that no player has any incentive to unilaterally change his strategy (see Nash (1950)). It is proven that if the players can randomize on their pure strategies set according to *any* probability distribution, then the game has at least one Nash equilibrium in the randomized mixed strategies. However, assume that Alice has restrictions on implementing probability distributions when choosing his mixed strategy. More specifically, assume that Alice can only choose a mixed action with entropy at most a given number h. For instance, if $\mathsf{h} = 0$, Alice is only allowed to play pure actions. We are interested in the maximum payoff that Alice can secure with mixed actions of entropy at most h (regardless of the action of Bob). Let us denote this by $\mathcal{J}(\mathsf{h})$. Equivalently, we may ask: in order to guarantee average payoff of w, how much randomness needs to be utilized by Alice? Let us denote the minimum entropy of the randomness consumed by Alice to guarantee payoff w by $F(w)$. We may call $F(w)$, the *min-entropy function*. [2] Observe that $F(\cdot)$ is the inverse function of $\mathcal{J}(\cdot)$.

Even though $\mathcal{J}(\mathsf{h})$ is defined in a one-shot setting, the (upper concave envelope) of $\mathcal{J}(\mathsf{h})$ finds an operational meaning in Gossner and Vieille (2002); Neyman and Okada (2000, 2009) in the context of repeated games as the payoff that Alice can secure in the long run. For instance, Gossner and Vieille (2002) study a repeated game in which Alice cannot randomize freely, but has access to an external i.i.d. source $X_1, X_2, \cdots$ which are revealed symbol by symbol (causally) to Alice as the game is played out. They show that the entropy of the source $H(X)$ characterizes the payoff that Alice can secure in the long run, and the answer is given in terms of $\mathcal{J}(H(X))$. This is similar to Shannon's compression formula $H(X)$ which is defined on a single copy of the source (*single letter*), but gives the ultimate compression limit when multiple copies of the source are observed.

[2] One should not confuse our "min-entropy function" with the term "min-entropy" commonly used to denote the Rényi entropy of order infinity.

As a generalization of the work of Gossner and Vieille (2002), we consider the case that Alice's observation sequence $X_1, X_2, \ldots$ is partially leaked to Bob. We show that even in this setting the security level in the long run is characterized by the function $\mathcal{J}(\cdot)$. Another motivation for $\mathcal{J}(\mathsf{h})$ is given in Hubek et al. (2016); Budinich and Fortnow (2011), arguing that limitation on random strategies of players stems from simplicity of humans; that "humans are known to be bad at generating random-like sequences".

To compute $F(w)$, one has to first consider the set of distributions on the action of Alice that would secure a payoff w for her. This set will be a polytope in the space of all probability distributions. Then, one should solve an entropy minimization problem over this polytope in the space of probability distributions. In fact, minimizing and maximizing entropy arises in a wide range of contexts. Computing maximum entropy under a set of linear constraints is a well-studied problem with a wide range of applications, *e.g.,* see Fang et al. (2012)(Cover and Thomas, 2012, p.367) and the principle of maximum entropy. It is shown in Kapur et al. (1995) that computing the minimum entropy can be also quite important. Furthermore, Watanabe (1981) shows that many algorithms for clustering and pattern recognition are essentially solving entropy minimization problems. An important special case of the entropy minimization problem (with its own applications) is that of finding a joint probability of minimum entropy given its marginal distributions (the marginal distribution is a linear constraint on the joint probability distribution); see Kocaoglu et al. (2016); Cicalese et al. (2017); Kovaevi et al. (2015). Finally, Shor (2004) shows that the quantum version of the entropy minimization problem is closely related to a number of important problems in quantum information theory.

It is shown by Kovacevic et al. (2012) that entropy minimization problem is an NP-hard non-convex optimization problem. Since the entropy is a concave function over the probability simplex, its minimum occurs at a vertex of the feasible domain. As a result, computation of $F(w)$ leads to a search problem over an exponentially large set. Even though the algorithm of Srikanth et al. (2003) can be used to compute $F(w)$, but it has no guarantee of success.

Our contributions can be summarized as follows:

- The min-entropy function $F(w)$ finds its operational interpretation in repeated games (e.g., the result of Gossner and Vieille (2002)). We simplify and generalize the proof of Gossner and Vieille (2002) in Sections 3 and 4. The generalized setup considers the possibility of leackage of Alice's random source sequence to Bob in the zero-sum game. More specifically, we assume an i.i.d. sequence of pairs (X_1, Y_1), $(X_2, Y_2), \cdots$ distributed according to a given $p(x, y)$. The sequence $X_1, X_2, \cdots$ is revealed symbol by symbol (causally) to Alice as the game is played out, while the sequence $Y_1, Y_2, \cdots$ is revealed symbol by symbol to Bob. We can view Y_i as the leackage that Bob obtains about Alice's observation. As before, Alice cannot randomize freely and can only use the randomness in the sequence $X_1, X_2, \cdots$. We show that the upper concave envelope of $\mathcal{J}(\mathsf{h})$, at $\mathsf{h} = H(X|Y)$ is the maximum payoff that Alice can secure regardless of Bob's actions. Furthermore, Bob can ensure that Alice does not get more than this payoff. Here $H(X|Y)$ is the conditional entropy of X given Y. This result is a generalization of the result of Gossner and Vieille since when Y is a constant random variable, the conditional entropy $H(X|Y)$ reduces to the unconditional entropy $H(X)$. We also show that if a genie provides the values of $Y_1, Y_2, \ldots$ symbol by symbol to Alice, the maximum payoff she can secure remains unchanged. In other words, knowledge of what Bob knows about Alice's observations is not helpful to Alice.

- We study the properties of $F(w)$ and give a number of easy-to-compute bounds on the value of $F(w)$. One of the bounds on $F(w)$ is expressed in terms of a linear program; this linear program has a quite general form. We employ *probabilistic tools* to study this linear program to prove the following explicit lower bound on $F(w)$:

$$F(w) \geq \log_2 \left(1 + \frac{(w-v)^2}{(w-\underline{m})(\overline{m}-w)}\right) \quad \forall w : v \leq w \leq w^*,$$

  where $\underline{m}$ and $\overline{m}$ are the minimum and maximum entries of the entries of the payoff table; v is the payoff that Alice can guarantee with deterministic strategies (pure-

strategy security level) and w^* is the Nash value of the game for Alice (payoff that Alice can guarantee with no restriction on her action's entropy). Because the linear program we consider has a quite general form (and can be converted to other linear programs with a change of variables), our use of probabilistic tools for studying linear programs is novel.

- We give and utilize new results on the set of mixed strategies that guarantee a certain security level. While the literature on game theory makes extensive use of total variation distance between distributions, we propose and illustrate the use of the χ^2–divergence (or the Tsallis divergence of order two) in game theoretic contexts. The χ^2–divergence is defined as $\chi^2(\mathbf{p},\mathbf{q}) = \left(\sum_i \frac{p_i^2}{q_i}\right) - 1$, where $\mathbf{p} = (p_1, p_2, \cdots, p_n)$ and $\mathbf{q} = (q_1, q_2, \cdots, q_n)$ are two probability distributions. To see the applicability of this divergence measure, consider a zero-sum game between Alice and Bob. Fix an arbitrary (mixed) strategy for Bob. Let $\mathbf{p}_1$ and $\mathbf{p}_2$ be two probability distributions for the action of Alice. Let W_1 and W_2 denote the payoff of Alice when she plays according to $\mathbf{p}_1$ and $\mathbf{p}_2$ respectively. Then, $\mathbb{E}[W_1]$ and $\mathbb{E}[W_2]$ are the payoffs secured by $\mathbf{p}_1$ and $\mathbf{p}_2$. Also $\mathsf{Var}[W_1]$ is a natural object, the risk associated to $\mathbf{p}_1$. Then, as we see later in Lemma 24, these natural objects in a game come together as follows:

$$\chi^2(\mathbf{p}_1, \mathbf{p}_2) \geq \frac{(\mathbb{E}[W_1] - \mathbb{E}[W_2])^2}{\mathsf{Var}[W_1]}.$$

The rest of this paper is organized as follows: In Section 2 we discuss our notation and provide a brief introduction to entropy of random variables; in Section 3 we simplify the proof of Gossner and Vieille (2002) and generalize it in Section 4. In Section 5 we study the computational aspects of the min-entropy function. Finally, in Section 6 we give the proofs. Some of the details are left for the appendices.

2. Preliminaries

2.1. Notation

We use the notation x_i^j to denote the sequence $(x_i, x_{i+1}, \ldots, x_j)$. When $i = 1$ we drop the subscript, and write x^j for $(x_1, x_2, \ldots, x_j)$. The same notation is used to represent sequences of random variables, *i.e.,* $X_i^j = (X_i, X_{i+1}, \ldots, X_j)$. Calligraphic letters such as $\mathcal{X}, \mathcal{Y}, \mathcal{A}, \mathcal{B}$ represent finite sets. Real vectors are represented by bold lower case letters. For example a probability mass function on a finite space $\mathcal{X} = \{1, 2, \ldots, n\}$, beside the representation $p(x)$, will be also denoted by a vector $\mathbf{p} = (p_1, \ldots, p_n)$.

We say that X^n is drawn i.i.d. from $p(x)$ if

$$p(x^n) = \prod_{i=1}^{n} p(x_i).$$

The total variation distance is denoted by $d_1(\cdot, \cdot)$ or $\|\cdot\|_{TV}$ and is defined as follows:

$$d_1(\mathbf{p}, \mathbf{q}) = \|\mathbf{p} - \mathbf{q}\|_{TV} \triangleq \frac{1}{2} \sum_i |p_i - q_i|.$$

2.2. Entropy function

Let $X \in \mathcal{X} = \{1, 2, \ldots, n\}$ be a random variable with probability mass function (pmf) $\mathbf{p} = (p_1, \ldots, p_n)$. The entropy of a random variable X quantifies the amount of randomness in X or the amount of information communicated by revealing the value of X and is defined to be:

$$H(X) = \sum_{i=1}^{n} -p_i \log(p_i)$$

where $0 \log(0) = 0$ by convention and all logarithms in this paper are in base two. As $H(X)$ is just a function of the distribution $\mathbf{p}$ we sometimes use notation $H(\mathbf{p})$ instead of $H(X)$. For a sequence of random variables $X^n = (X_1, X_2, \ldots, X_n) \in \mathcal{X}^n$ with joint pmf $p(x^n)$, we have

$$H(X^n) = \sum_{x^n \in \mathcal{X}^n} -p(x^n) \log(p(x^n)).$$

For random variable pair $(X, Y) \in \mathcal{X} \times \mathcal{Y}$ with pmf $p(x, y)$, the conditional entropy of X given Y is the amount of average uncertainty about X for someone who already knows the value of Y, and is defined as follows:

$$H(X|Y) = \sum_{(x,y)\in\mathcal{X}\times\mathcal{Y}} -p(x, y) \log(p(x|y)).$$

Using the definition of entropy and conditional entropy one can check that

$$H(X, Y) = H(Y) + H(X|Y)$$

holds. Utilizing the above property iteratively, one obtains the chain rule for a sequence of random variables:

$$H(X^n) = H(X_1) + \sum_{i=2}^{n} H(X_i|X^{i-1}).$$

Following properties hold for the entropy function:

- $H(X) \geq 0$.
- $H(X_1, X_2) \geq H(X_1)$.
- Let $f(x)$ be arbitrary deterministic function then $H(f(X)) \leq H(X)$.

3. Simplifying the proof of Gossner and Vieille (2002) for repeated games with a bounded randomness source

In this section we revisit the repeated game studied in Gossner and Vieille (2002). While the high level picture of the proof of Gossner and Vieille (2002) is simple: it uses the so-called "block-Markov" proof technique of information theory, where we divide time into a number of blocks; Alice uses its observation in each block to produce its actions for the next block. To make their argument precise, Gossner and Vieille (2002) give an elaborate proof (see Proposition 8 and its proof in Appendix A). In particular, Gossner and Vieille (2002) introduce a new notion of *absolute Kullback distance* which is not used in the information theory literature. Our goal is to show the high level picture of the proof of Gossner and Vieille (2002) can be made precise in an easier manner using standard information theory tools.

3.1. Problem definition

Let us first begin with reviewing the definition of the repeated game studied in Gossner and Vieille (2002). Consider an T stage repeated zero-sum game between players Alice(A) and Bob(B) with respective finite action sets $\mathcal{A} = \{1, \ldots, n\}$ and $\mathcal{B} = \{1, \ldots, n'\}$ where n and n' are natural numbers. Let $X^T = (X_1, X_2, \ldots, X_T)$ be a sequence of random variables drawn i.i.d. from sample space $\mathcal{X}$ with law p_X. In every stage $t \in \{1, 2, \ldots, T\}$ first Alice observe random source X_t privately then Alice and Bob choose actions $A_t \in \mathcal{A}$ and $B_t \in \mathcal{B}$ respectively. At the end of stage t both players observe the chosen actions A_t and B_t and Alice gets stage payoff u_{A_t,B_t} from Bob. In order to choose actions at stage t, players make use of the history of their observations until stage t which is denoted by $\mathsf{H}_1^t = (X_1, A_1, B_1, \ldots, X_{t-1}, A_{t-1}, B_{t-1}, X_t)$ for Alice and $\mathsf{H}_2^t = (A_1, B_1, \ldots, A_{t-1}, B_{t-1})$ for Bob. Let $\sigma_t : (\mathcal{A} \times \mathcal{B})^{t-1} \times \mathcal{X}^t \to \mathcal{A}$ and $\tau_t : (\mathcal{A} \times \mathcal{B})^{t-1} \to \mathcal{B}$ be the functions mapping the history of observations of Alice and Bob to actions at stage t, thus $A_t = \sigma_t(\mathsf{H}_1^t)$ and $B_t = \tau_t(\mathsf{H}_2^t)$. Note that Alice and Bob do not have access to any private sources of randomness besides H_1^t and H_2^t, and have to use deterministic functions $\sigma_t(\cdot)$ and $\tau_t(\cdot)$. We call the T-tuples $\sigma = (\sigma_1, \sigma_2, \ldots, \sigma_T)$ and $\tau = (\tau_1, \tau_2, \ldots, \tau_T)$ the strategies of Alice and Bob respectively. The expected average payoff for Alice up to stage T induced by strategies σ and τ is denoted by $\lambda_T(\sigma, \tau)$ which is:

$$\lambda_T(\sigma, \tau) = \mathbb{E}_{\sigma,\tau}\left[\frac{1}{T}\sum_{t=1}^{T} u_{A_t,B_t}\right]$$

where $\mathbb{E}_{\sigma,\tau}$ denotes expectation with respect to the distribution induced by i.i.d. repetitions of p_X and strategies σ and τ. Alice (Bob) wishes to maximize (minimize) $\lambda_T(\sigma, \tau)$.

Definition 1. *Let v be arbitrary real value:*

- *Alice can secure v if there exists a strategy σ^* for Alice such that for all strategy τ of Bob we have* $\liminf_{T\to\infty} \lambda_T(\sigma^*, \tau) \geq v$.
- *Bob defends v if given an arbitrary strategy σ for Alice there exists a strategy τ^* for Bob such that* $\limsup_{T\to\infty} \lambda_T(\sigma, \tau^*) \leq v$.

- *v is the maxmin value of the repeated game if Alice can secure v and Bob can defend v.*

3.2. Proof of Gossner and Vieille (2002) for the payoff that Alice can secure

As mentioned in Gossner and Vieille (2002), the upper concave envelope $\mathcal{J}_{\text{cav}}(\cdot)$ at $\mathsf{h} = H(X)$ can be expressed as the convex combination

$$\gamma J(H(p_A^{(1)})) + (1-\gamma)J(H(p_A^{(2)}))$$

for some $\gamma \in [0,1]$ and input distributions $p_A^{(1)}$ and $p_A^{(2)}$ on $\mathcal{A}$ that secure payoffs $\mathcal{J}(H(p_A^{(1)}))$ and $\mathcal{J}(H(p_A^{(2)}))$ respectively in the one-shot game, and satisfy

$$\gamma H(p_A^{(1)}) + (1-\gamma)H(p_A^{(2)}) = H(X).$$

As a result, it suffices to show that for any $p_A^{(1)}$ and $p_A^{(2)}$ satisfying

$$\gamma H(p_A^{(1)}) + (1-\gamma)H(p_A^{(2)}) < H(X),$$

Alice can use the source of randomness X to secure payoffs arbitrarily close to

$$\lambda \min_{b\in\mathcal{B}} \mathbb{E}_{p_A^{(1)}}[u_{A,b}] + (1-\lambda) \min_{b\in\mathcal{B}} \mathbb{E}_{p_A^{(2)}}[u_{A,b}],$$

which is the weighted sum of the payoffs that input distributions $p_A^{(1)}$ and $p_A^{(2)}$ secure.

The game is played for T stages. Take some T of the form $T = NL$, and divide the total T stages into N blocks of length L. Excluding the first block, we generate the action sequence of each block as a function of the random source observed during the previous block (we ignore the payoff of the first block throughout the discussion, since by taking the number of blocks N large enough, the contribution of the first block in Alice's net average payoff becomes negligible). More specifically, excluding the first block, each block is further divided into two subblocks, with the first subblock taking up γ fraction of the block as in Fig. 1. Alice aims to use its observed random source during the previous block to play almost i.i.d. according to $p_A^{(1)}$ during the first subblock, and almost i.i.d. according to $p_A^{(2)}$ during the second subblock. In addition, Alice wants its action at any

given stage to be also (almost) independent of Bob's previous observation history in that stage. Observe that if Alice could produce actions that are perfectly i.i.d. according to $p_A^{(1)}$ and independent of Bob's history, then in γ fraction of the games in a block, she secures the average payoff of $\min_{b\in\mathcal{B}} \mathbb{E}_{p_A^{(1)}}[u_{A,b}]$ per action, and in the remaining $1-\gamma$ fraction secures the payoff of $\min_{b\in\mathcal{B}} \mathbb{E}_{p_A^{(2)}}[u_{A,b}]$. This gives Alice a total payoff of

$$\lambda \min_{b\in\mathcal{B}} \mathbb{E}_{p_A^{(1)}}[u_{A,b}] + (1-\lambda) \min_{b\in\mathcal{B}} \mathbb{E}_{p_A^{(2)}}[u_{A,b}],$$

and will complete the proof.

It remains to show that Alice can play almost i.i.d. according to $p_A^{(1)}$ during the first subblock of each block, and almost i.i.d. according to $p_A^{(2)}$ during the second subblock of each block. Remember that $\gamma H(p_A^{(1)}) + (1-\gamma) H(p_A^{(2)}) < H(X)$. Intuitively speaking, since Alice's observation from a block has entropy $LH(X)$, which is larger than the entropy of its intended action in the next block $L\gamma H(p_A^{(1)}) + L(1-\gamma)H(p_A^{(2)})$, she should be able to find a proper mapping to produce her actions in the current block from her observations in the preceding block. Gossner and Vieille (2002) propose such a mapping and put effort to prove its correctness. This is where most of the effort is spent. The essential problem here is when a source can be used to simulate a source from another source (here the observations from one block to actions in the next block). This is problem is solved in the information theory literature. As stated below in Lemma 2 which is adopted from (Han, 2003, p. 110), the measure of relevance to compare between the two sources is *inf-entropy* and *sup-entropy*, rather than entropy. However, for i.i.d. sources (or concatenation of i.i.d. sources), inf-entropy and sup-entropy reduce to the normal Shannon entropy.

Lemma 2 (Simulation of a source from another source)**.** *For each natural number L, consider an arbitrary distribution $p_{X^L}^{(L)}$ on sequences $(x_1, x_2, \cdots, x_L) \in \mathcal{X}^L$. Similarly, for each natural number L, consider an arbitrary distribution $q_{A^L}^{(L)}$ on sequences $(a_1, a_2, \cdots, a_L) \in \mathcal{A}^L$. Let sup-entropy and inf-entropy of these sources be defined as*

$$\underline{H}(X) = p\text{-}\liminf_{L\to\infty} \frac{1}{L} \log \frac{1}{p_{X^L}(X^L)},$$

$$\overline{H}(A) = p\text{-}\limsup_{L\to\infty} \frac{1}{L} \log \frac{1}{p_{A^L}(A^L)},$$

where for a random sequence $\{Z_t\}$,

$$p\text{-}\liminf_{t\to\infty} Z_t = \sup\left\{\beta | \lim_{t\to\infty} \Pr[Z_t < \beta] = 0\right\},$$

$$p\text{-}\limsup_{t\to\infty} Z_t = \inf\left\{\beta | \lim_{t\to\infty} \Pr[Z_t > \beta] = 0\right\}.$$

Note that if $p_{X^L}^{(L)}$ *and* $p_{A^L}^{(L)}$ *are i.i.d. distributions according to some* $p(x)$ *and* $p(a)$ *respectively, then* $\underline{H}(X) = \overline{H}(X) = H(X)$ *and* $\underline{H}(A) = \overline{H}(A) = H(A)$.

Then, if $\underline{H}(X) > \overline{H}(A)$, *for each natural number* L, *there exists a mapping* $\varphi_L : \mathcal{X}^L \to \mathcal{A}^L$ *such that the total variation distance between the distributions of* $\varphi_L(X^L)$ *and* A^L *vanishes asymptotically,* i.e.,

$$\lim_{L\to\infty} \|q_{A^L}^{(L)} - p_{\varphi_L(X^L)}^{(L)}\|_{TV} = 0,$$

where $\|\cdot\|_{TV}$ *is the total variation distance.*

Remark 3. *Observe that* $\lim_{L\to\infty} \|p_{A^L} - p_{\varphi_L(X^L)}\|_{TV} = 0$ *in the above lemma implies that* $\varphi_L(X^L)$ *is almost statistically indistinguishable from* p_{A^L} *for large values of* L. *In other words, there is* <u>no statistical test</u> *that can distinguish between* p_{A^L} *and* $p_{\varphi_L(X^L)}$ *with a non-negligible probability.*

Take some $\epsilon > 0$. Since $\lambda H(p_A^{(1)}) + (1-\lambda)H(p_A^{(2)}) < H(X)$ from Lemma 2 there exists a mapping $\varphi : \mathcal{X}^L \to \mathcal{A}^L$ such that for large L, the pmf of $A^L = \varphi(X^L)$ is approximated (in total variation distance) as

$$\left\| p_{A^L}(a^L) - \prod_{t=1}^{\lceil \gamma L \rceil} p_A^{(1)}(a_t) \prod_{t=\lceil \gamma L \rceil + 1}^{L} p_A^{(2)}(a_t) \right\|_{TV} \leq \epsilon. \tag{1}$$

Therefore by dividing the T stages into N blocks of L stages, we construct a strategy σ for Alice as follows: in each block (excluding the first block) Alice chooses action sequence $A^L = \varphi(X^L)$ where X^L is the randomness source observed during the previous block; *i.e.*, A^L is for the current block but X^L is for the previous block.

The ideal distribution $\prod_{t=1}^{\lceil \gamma L \rceil} p_A^{(1)}(a_t) \prod_{t=\lceil \gamma L \rceil+1}^{L} p_A^{(2)}(a_t)$ gives Alice a payoff of

$$\lceil \gamma L \rceil \min_{b \in \mathcal{B}} \mathbb{E}_{p_A^{(1)}}[u_{A,b}] + (L - \lceil \gamma L \rceil) \min_{b \in \mathcal{B}} \mathbb{E}_{p_A^{(2)}}[u_{A,b}].$$

Alice's actual distribution is within ϵ total variation distance of the ideal distribution. By relating the total variation distance to the payoff differences, we obtain that difference between the payoff under the actual distribution and the ideal one is at most ϵ times $2L\mathsf{M}$ where M is the maximum absolute value entry of the game table ($L\mathsf{M}$ is the maximum absolute value entry of the L repetitions of the game). Thus, the distance between the average payoff of the block differs by at most $2\epsilon\mathsf{M}$ from the average payoff under the ideal distribution. This completes the proof.

Remark 4. *We constructed the strategy σ for the case $T = NL$. In general for $T = NL+\delta$ where $\delta < L$ one can extend the first block to contain $L + \delta$ stages and choose N large enough to diminish the effect of first block on average payoff.*

4. Extending the proof of Gossner and Vieille (2002) to include leakage of the randomness source

We generalize the model of repeated game in Section 3 as follows: Let $X^T = (X_1, X_2, \ldots, X_T)$ and $Y^T = (Y_1, Y_2, \ldots, Y_T)$ be two sequences of random variables, such that (X_t, Y_t) for $t \in \{1, 2, \ldots, T\}$, are drawn i.i.d. from a given pmf p_{XY}. In every stage $t \in \{1, 2, \ldots, T\}$, Alice and Bob respectively observe random sources X_t and Y_t privately, and choose actions $A_t \in \mathcal{A}$ and $B_t \in \mathcal{B}$. Thus the only difference of this model with the model of Section 3 is that in every stage t Bob observes Y_t which is related to the observation of Alice X_t, hence the history of observations of Bob is $\mathsf{H}_2^t = (Y_1, A_1, B_1, \ldots, Y_{t-1}, A_{t-1}, B_{t-1}, Y_t)$ and $B_t = \tau_t(\mathsf{H}_2^t)$ where $\tau_t : (\mathcal{A} \times \mathcal{B})^{t-1} \times \mathcal{Y}^t \to \mathcal{B}$; on the other hand, the history of observations of Alice remains unchanged. Note that Alice and Bob do not have access to any private source of randomness besides H_1^t and H_2^t, and have to use deterministic functions $\sigma_t(\cdot)$ and $\tau_t(\cdot)$. Alice (Bob) wishes to maximize (minimize) the expected average payoff $\lambda_T(\sigma, \tau)$.

The above repeated zero-sum game between players Alice(A) and Bob(B) reduces to the one considered in Section 3 if we set Y_t to be constant random variables. The maxmin value of the repeated game under study is characterized as follows:

Theorem 5. *The maxmin value of the above repeated game is* $\mathcal{J}_{cav}(H(X|Y))$.

Corollary 6. *Consider the case in which besides the sequence* $X_1, X_2, \ldots$, *Alice also observes the sequence* $Y_1, Y_2, \ldots$, *therefore* $X'_t = (X_t, Y_t)$ *is the observation of Alice at stage* t. *Since* $H(X'|Y) = H(X|Y)$, *the maxmin value of the game is still* $\mathcal{J}_{cav}(H(X|Y))$ *which is equal to the maxmin value of the game when Alice does not observe the sequence* $Y_1, Y_2, \ldots$. *Therefore, Alice's access to* $Y_1, Y_2, \ldots$ *does not help her. This parallels the classical result of Slepian and Wolf in information theory Slepian and Wolf (1973).*

Example 7. *(Matching pennies) Consider the matching pennies game with* $\mathcal{A} = \mathcal{B} = \{0, 1\}$ *and payoffs:*

$$u_{0,0} = u_{1,1} = 1, \quad u_{0,1} = u_{1,0} = 0.$$

Assume that Alice observes (symbol by symbol) an i.i.d. sequence of binary random variables $X_1, X_2, \cdots$, *where*

$$p_{X_i}(0) = q, \quad p_{X_i}(1) = 1 - q.$$

Each X_i *is revealed to Bob with probability* $\alpha \in [0, 1]$. *Thus, Bob observes (symbol by symbol) a sequence* $Y_1, Y_2, \cdots$ *such that with probability* α, $Y_i = X_i$ *and with probability* $1 - \alpha$, $Y_i = null$.

Roughly speaking α *fraction of the bits of Alice are revealed to Bob. If Alice is made aware of which of its bits are compromised and leaked to Bob, she can drop the compromised bits and keep the* $1 - \alpha$ *fraction of the secret bits. In this way, she can distill an average of* $(1 - \alpha)H(q)$ *random bits per observation. It is known from Gossner and Vieille (2002) that for the matching pennies game* $\mathcal{J}_{cav}(\mathsf{h}) = \mathsf{h}/2$. *Thus, using* $(1-\alpha)H(q)$ *bits randomness per stage, Alice can secure a payoff* $(1 - \alpha)H(q)/2$. *The above theorem says that Alice can secure the same payoff even without knowing which of its bits are compromised.*

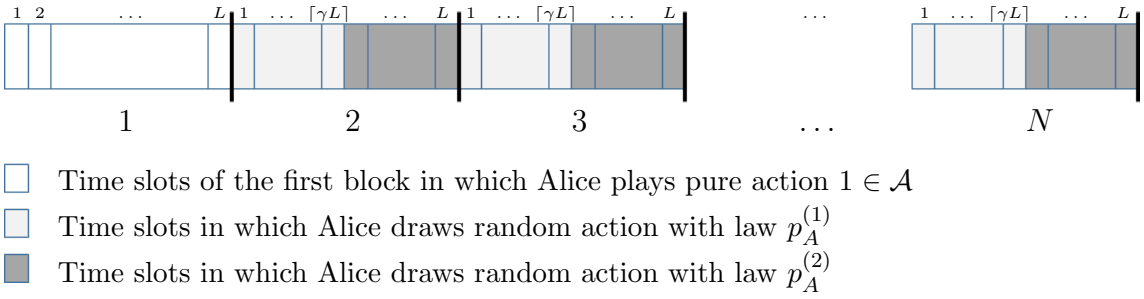
1 2 ... L 1 ... ⌈γL⌉ ... L 1 ... ⌈γL⌉ ... L ... 1 ... ⌈γL⌉ ... L
1
2
3
...
N
Time slots of the first block in which Alice plays pure action 1 ∈ 𝒜
Time slots in which Alice draws random action with law $p_A^{(1)}$
Time slots in which Alice draws random action with law $p_A^{(2)}$

Figure 1: Illustration of block Markov strategy

The proof uses a known fact in information theory that $H(X|Y)$ is the amount of random bits extracted from source $X_1, X_2, \dots$ that are almost independent of the observations of Bob $(Y_1, Y_2, \dots)$. As in Section 3, the optimal strategy for Alice utilizes the random sources to extract randomness from the source and use them to randomize her actions. To prove Theorem 5, in Section 4.1 we show that Alice can secure $\mathcal{J}_{\text{cav}}(H(X|Y))$ and in Section 4.2 we show that Bob can defend $\mathcal{J}_{\text{cav}}(H(X|Y))$. Therefore, by definition $\mathcal{J}_{\text{cav}}(H(X|Y))$ is the maxmin value of the repeated game.

4.1. Alice can secure $\mathcal{J}_{cav}(H(X|Y))$

To show that Alice can secure any payoff less than $\mathcal{J}_{\text{cav}}(H(X|Y))$, we extend and simplify the proof of Gossner and Vieille (2002). As before, the upper concave envelope $\mathcal{J}_{\text{cav}}(\cdot)$ at $\mathsf{h} = H(X|Y)$ can be expressed as the convex combination

$$\gamma J(H(p_A^{(1)})) + (1-\gamma) J(H(p_A^{(2)}))$$

for some $\gamma \in [0,1]$ and input distributions $p_A^{(1)}$ and $p_A^{(2)}$ on $\mathcal{A}$ that secure payoffs $\mathcal{J}(H(p_A^{(1)}))$ and $\mathcal{J}(H(p_A^{(2)}))$ respectively in the one-shot game, and satisfy

$$\gamma H(p_A^{(1)}) + (1-\gamma) H(p_A^{(2)}) = H(X|Y).$$

As a result, it suffices to show that for any $p_A^{(1)}$ and $p_A^{(2)}$ satisfying

$$\gamma H(p_A^{(1)}) + (1-\gamma) H(p_A^{(2)}) < H(X|Y),$$

Alice can secure payoffs arbitrarily close to

$$\lambda \min_{b\in\mathcal{B}} \mathbb{E}_{p_A^{(1)}}[u_{A,b}] + (1-\lambda) \min_{b\in\mathcal{B}} \mathbb{E}_{p_A^{(2)}}[u_{A,b}].$$

Again, the idea is to utilize the block-Markov proof technique. Take some T of the form $T = NL$, and divide the total T stages into N blocks of length L. Each block is divided into two subblocks as before. Excluding the first block, Alice wants to play almost i.i.d. according to $p_A^{(1)}$ during the first subblock, and almost i.i.d. according to $p_A^{(2)}$ during the second subblock.

By symmetry, we only need to consider the payoff that Alice gets in one of the blocks. For notational simplicity, we denote the observations of Alice and Bob in the *previous* block by X^L and Y^L respectively, and use A^L and B^L to denote their actions in the *current* block. Instead of Lemma 2 in the previous section, the proof relies on the following proposition whose proof is given in Section 6.1:

Proposition 8. *Let $0 \le \gamma \le 1$ be an arbitrary real number and $p_A^{(1)}$ and $p_A^{(2)}$ be arbitrary distributions on $\mathcal{A}$ such that $\gamma H(p_A^{(1)}) + (1-\gamma)H(p_A^{(2)}) < H(X|Y)$. Then, for any $\epsilon > 0$, there exists a natural number L and a mapping $\psi : \mathcal{X}^L \to \mathcal{A}^L$ such that*

$$\left\| p_{A^L,Y^L}(a^L, y^L) - p_{Y^L}(y^L) \prod_{t=1}^{\lceil \gamma L \rceil} p_A^{(1)}(a_t) \prod_{t=\lceil \gamma L \rceil + 1}^{L} p_A^{(2)}(a_t) \right\|_{TV} < \epsilon, \tag{2}$$

where (X^L, Y^L) are i.i.d. according to $p(x,y)$ and $A^L = \psi(X^L)$.

Note that $p_{Y^L}(y^L) \prod_{t=1}^{\lceil \gamma L \rceil} p_A^{(1)}(a_t) \prod_{t=\lceil \gamma L \rceil+1}^{L} p_A^{(2)}(a_t)$ is the ideal independent actions by Alice, whereas $p_{A^L,Y^L}(a^L, y^L)$ is the real joint distribution.

Let Alice use the actions $(A_1, A_2, \cdots, A_L)$ in the current block where $A^L = \psi(X^L)$. Bob's action at time t is a deterministic function of the history of actions and his observations at the previous Y^L, and his observations in the current block. Thus

$$p(b^L | y^L, a^L) = \prod_{t=1}^{L} p(b_t | y^L, a^{t-1}, b^{t-1})$$

Then, utilizing the following property of total variation for random variables Y and Z

$$\|p_Y p_{Z|Y} - q_Y p_{Z|Y}\|_{TV} = \|p_Y - q_Y\|_{TV},$$

we conclude from (2) that

$$\Big\| p_{A^L,Y^L}(a^L, y^L) \prod_{t=1}^{L} p(b_t|y^L, a^{t-1}, b^{t-1}) \tag{3}$$
$$- p_{Y^L}(y^L) \prod_{t=1}^{\lceil \gamma L \rceil} p_A^{(1)}(a_t) p(b_t|y^L, a^{t-1}, b^{t-1}) \prod_{t=\lceil \gamma L \rceil + 1}^{L} p_A^{(2)}(a_t) p(b_t|y^L, a^{t-1}, b^{t-1}) \Big\|_{TV} < \epsilon,$$

Utilizing the following property of total variation for random variables Y and Z

$$\|p_{Y,Z} - q_{Y,Z}\|_{TV} \geq \|p_Y - q_Y\|_{TV},$$

we conclude that the distance between the actual actions $p_{A^L,B^L}(a^L, b^L)$ and the ideal one is less than or equal to ϵ, *i.e.,*

$$\Big\| p_{A^L,B^L}(a^L, b^L) - \prod_{t=1}^{\lceil \gamma L \rceil} p_A^{(1)}(a_t) p(b_t|a^{t-1}, b^{t-1}) \prod_{t=\lceil \gamma L \rceil + 1}^{L} p_A^{(2)}(a_t) p(b_t|a^{t-1}, b^{t-1}) \Big\|_{TV} < \epsilon. \tag{4}$$

As before, by relating the total variation distance to the payoff differences, we obtain that the average payoff differs by at most $2\epsilon \mathsf{M}$ from the average payoff under the ideal distribution. This completes the proof.

4.2. Bob can defend $\mathcal{J}_{cav}(H(X|Y))$

This is an extension of the proof given in Neyman and Okada (2000); Gossner and Vieille (2002). Let σ be an arbitrary strategy for Alice and generate strategy τ for Bob as follows: given h_2^t, an arbitrary history of observations of Bob until stage t, $\tau_t(h_2^t)$ is the best choice of Bob that minimizes the expected payoff at stage t:

$$\tau_t(h_2^t) \in \arg\min_{j \in \mathcal{B}} \mathbb{E}_\sigma \left[u_{A_t,j} | \mathsf{H}_2^t = h_2^t \right]$$

where $\mathbb{E}_\sigma$ denotes expectation with respect to the probability distribution induced by σ and $p_{X,Y}$. Thus:

$$\mathbb{E}_{\sigma,\tau} \left[u_{A_t,B_t} | \mathsf{H}_2^t = h_2^t \right] \leq \mathcal{J} \left(H(A_t|\mathsf{H}_2^t = h_2^t) \right) \leq \mathcal{J}_{\text{cav}} \left(H(A_t|\mathsf{H}_2^t = h_2^t) \right), \tag{5}$$

since conditional on the observations h_2^t of Bob, Alice's action A_t has entropy $H(A_t|\mathsf{H}_2^t = h_2^t)$. Therefore:

$$\begin{aligned}
\mathbb{E}_{\sigma,\tau}[u_{A_t,B_t}] &= \sum_{h_2^t \in (\mathcal{A}\times\mathcal{B})^{t-1}\times\mathcal{X}_2^t} \Pr\left[\mathsf{H}_2^t = h_2^t\right] \mathbb{E}_{\sigma,\tau}\left[u_{A_t,B_t}|\mathsf{H}_2^t = h_2^t\right] \\
&\leq \sum_{h_2^t \in (\mathcal{A}\times\mathcal{B})^{t-1}\times\mathcal{X}_2^t} \Pr\left[\mathsf{H}_2^t = h_2^t\right] \mathcal{J}_{\text{cav}}\left(H(A_t|\mathsf{H}_2^t = h_2^t)\right) \\
&\leq \mathcal{J}_{\text{cav}}\left(\sum_{h_2^t \in (\mathcal{A}\times\mathcal{B})^{t-1}\times\mathcal{X}_2^t} \Pr\left[\mathsf{H}_2^t = h_2^t\right] H(A_t|\mathsf{H}_2^t = h_2^t)\right) \\
&= \mathcal{J}_{\text{cav}}\left(H(A_t|\mathsf{H}_2^t)\right). \qquad (6)
\end{aligned}$$

The first inequality results from (5) and by using Jensen's inequality for concave function $\mathcal{J}_{\text{cav}}(\cdot)$ the second inequality follows. By definition of $\lambda_T(\sigma,\tau)$ and using (6) we have:

$$\begin{aligned}
\lambda_T(\sigma,\tau) &\leq \sum_{t=1}^{T} \frac{1}{T} \mathcal{J}_{\text{cav}}\left(H(A_t|\mathsf{H}_2^t)\right) \\
&\leq \mathcal{J}_{\text{cav}}\left(\sum_{t=1}^{T} \frac{1}{T} H(A_t|\mathsf{H}_2^t)\right) \\
&= \mathcal{J}_{\text{cav}}\left(\sum_{t=1}^{T} \frac{1}{T} H(A_t|Y^t, A^{t-1}, B^{t-1})\right) \\
&= \mathcal{J}_{\text{cav}}\left(\sum_{t=1}^{T} \frac{1}{T} H(A_t|Y^T, A^{t-1})\right) \\
&= \mathcal{J}_{\text{cav}}\left(\frac{1}{T} H(A^T|Y^T)\right) \\
&\leq \mathcal{J}_{\text{cav}}\left(\frac{1}{T} H(X^T|Y^T)\right) \\
&= \mathcal{J}_{\text{cav}}\left(H(X|Y)\right) \qquad (7)
\end{aligned}$$

where the second inequality follows by using Jensen's inequality for concave function $\mathcal{J}_{\text{cav}}$. Note that given the strategy τ for Bob, B^{t-1} is a function of Y^t and A^{t-1} thus $H(A_t|Y^t, A^{t-1}, B^{t-1}) = H(A_t|Y^t, A^{t-1})$ and as A_t and B_t are causally generated from i.i.d. sequences X^T and Y^T hence (A_t, A^{t-1}, Y^t) is independent of $(X_{2,t+1}, X_{2,t+2}, \ldots, X_{2,T})$ and $H(A_t|Y^t, A^{t-1}) = H(A_t|Y^T, A^{t-1})$ follows. Therefore the second equality is correct. To

prove the final equality consider that given strategies of the players and the sequence Y^T, A^T is a function of X^T thus $H(X^T|Y^T) \geq H(A^T|Y^T)$ and as $\mathcal{J}_{\text{cav}}(\cdot)$ is an increasing function, the final equality holds.

Thus for all strategy σ of Alice there exists a strategy τ such that $\limsup_{T\to\infty} \lambda_T(\sigma,\tau) \leq \mathcal{J}_{\text{cav}}(H(X|Y))$. Hence Bob can defend $\mathcal{J}_{\text{cav}}(H(X|Y))$.

5. On the computation of min-entropy function

Consider a zero-sum game between players Alice (A) and Bob (B) with respective pure strategy sets $\mathcal{A} = \{1,\ldots,n\}$ and $\mathcal{B} = \{1,\ldots,n'\}$ where n and n' are natural numbers. The payoff matrix is denoted by $\mathtt{U} = [u_{ij}]$ where u_{ij} is the real valued payoff that player A gets from player B when $i \in \mathcal{A}$ and $j \in \mathcal{B}$ are played. Player A (B) wishes to maximize (minimize) the expected payoff. The set of all randomized strategies of players A and B are denoted by $\Delta(\mathcal{A})$ and $\Delta(\mathcal{B})$ respectively which are probability simplexes on sample spaces $\mathcal{A}$ and $\mathcal{B}$ respectively. Thus, Alice's strategy corresponds to a pmf $\mathbf{p} = (p_1, p_2, \ldots, p_n)$, which can be also illustrated as a column vector with non-negative entries that add up to one.

Assume that player A uses randomized strategy $\mathbf{p}$. Then the payoff of Alice if Bob plays $j \in \mathcal{B}$ is $\sum_i p_i u_{ij}$. We say that Alice secures payoff w with randomized strategy $\mathbf{p}$ (regardless of the action of player B) if $\sum_i p_i u_{ij} \geq w$ for all $j \in \mathcal{B}$. Thus, the set of all distributions that guarantee payoff w for player A can be expressed as

$$\mathcal{P}_{\mathtt{U}}(w) = \{\mathbf{p} \in \Delta(\mathcal{A}) : \mathbf{p}^T\mathtt{U} \geq w\mathbf{1}^T\}, \tag{8}$$

where $\mathbf{p}^T$ is the transpose of the column vector $\mathbf{p}$, $\mathbf{1}$ is a column vector of all ones and $\mathbf{v}_1 \geq \mathbf{v}_2$ means any element of $\mathbf{v}_1$ is greater than or equal to the corresponding element at $\mathbf{v}_2$.

We define

$$F_{\mathtt{U}}(w) \triangleq \min_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} H(\mathbf{p}), \tag{9}$$

If the set $\mathcal{P}_{\mathtt{U}}(w)$ is empty, we set $F_{\mathtt{U}}(w) = +\infty$.

A remark on notation: Two-player zero-sum games are completely characterized by payoff matrix. For the sake of simplicity, we call two-player zero-sum games with their payoff matrix. Thus, game $\mathtt{U}$ refers to a game with payoff table $\mathtt{U}$.

Definition 9. *Given a game* $\mathtt{U} = [u_{ij}]$*, parameters* $\underline{m}$*,* $\overline{m}$*,* v *and* w^* *are defined as:*

- $\underline{m}$ *(*$\overline{m}$*) is the minimum (maximum) element of matrix* $\mathtt{U}$ *:* $\underline{m} = \min_{i,j} u_{ij}$ *(*$\overline{m} = \max_{i,j} u_{ij}$*).*
- v *is the maximum payoff secured by pure actions (*pure-strategy security level*):* $v = \max_i \min_j u_{ij}$.
- w^* *is the Nash value of* $\mathtt{U}$*, which is the maximum guaranteed payoff with unlimited access to random sources:*

$$w^* = \max_{w:\ \mathcal{P}_{\mathtt{U}}(w) \neq \emptyset} w. \tag{10}$$

Note that by definition, $\underline{m} \leq v \leq w^* \leq \overline{m}$.

According to Definition 9, v is the payoff that is guaranteed without consumption of any randomness whereas w^* is the maximum guaranteed payoff when unlimited randomness is available. Thus, it is interesting to consider the min-entropy function $F_{\mathtt{U}}(w)$ in the domain $v \leq w \leq w^*$. If $w \leq v$, then $F_{\mathtt{U}}(w) = 0$; if $w > w^*$ the feasible set of optimization problem in 9 is empty and $F_{\mathtt{U}}(w) = +\infty$. When $v \leq w \leq w^*$, the function $F_{\mathtt{U}}(w)$ is not necessarily convex or concave as a function of w: it is strictly increasing and piecewise concave (Neyman and Okada, 2000, p. 241). [3]

5.1. On the set $\mathcal{P}_{\mathtt{U}}(w)$

The function $F_{\mathtt{U}}(w)$ is defined in (9) using $\mathcal{P}_{\mathtt{U}}(w)$, which is the set of all distributions that guarantee a security level w for player A. Observe that

$$\mathcal{P}_{\mathtt{U}}(w) = \{\mathbf{p} \in \Delta(\mathcal{A}) : \mathbf{p}^T \mathtt{U} \geq w\mathbf{1}^T\}, \tag{11}$$

[3] This property is stated in Neyman and Okada (2000) in terms of the function $J(\cdot)$.

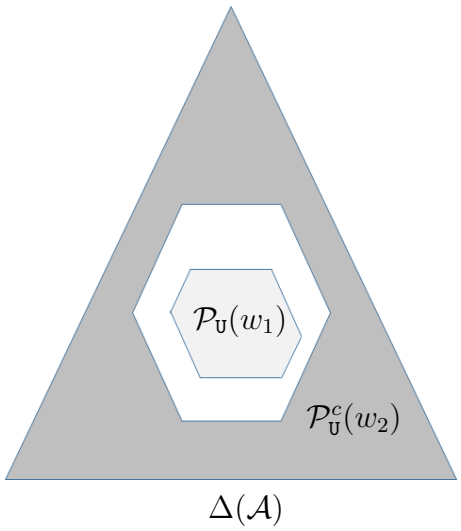

Figure 2: Illustration of $\mathcal{P}_{\mathtt{U}}(w_1)$, $\mathcal{P}^c_{\mathtt{U}}(w_2)$ and $\Delta(\mathcal{A})$

is a polytope defined via some linear constraints. As the matrix $\mathtt{U}$ is completely arbitrary, with a change of variables, one can convert it to many different equivalent linear programs. We only need to study $\mathcal{P}_{\mathtt{U}}(w)$ for $v \le w \le w^*$. It is immediate from the definition $\mathcal{P}_{\mathtt{U}}(w)$ that this set is decreasing in w, *i.e.,* for any $w_1 \ge w_2$,

$$\mathcal{P}_{\mathtt{U}}(w_1) \subseteq \mathcal{P}_{\mathtt{U}}(w_2). \tag{12}$$

We are interested to see if the inclusion in (12) is strict, and if yes, quantify to what extent it is. To do this, we look at the distance between the set $\mathcal{P}_{\mathtt{U}}(w_1)$ and the compliment of $\mathcal{P}_{\mathtt{U}}(w_2)$ (that is $\mathcal{P}^c_{\mathtt{U}}(w_2) = \Delta(\mathcal{A}) - \mathcal{P}_{\mathtt{U}}(w_2)$). The sets $\mathcal{P}_{\mathtt{U}}(w_1)$, $\mathcal{P}^c_{\mathtt{U}}(w_2)$ and $\Delta(\mathcal{A})$ are illustrated in Figure 2. The distance between any two sets can be defined as

$$d(\mathcal{S}_1, \mathcal{S}_2) \triangleq \inf_{\mathbf{p}\in\mathcal{S}_1, \mathbf{q}\in\mathcal{S}_2} d(\mathbf{p}, \mathbf{q}),$$

where $d(\mathbf{p}, \mathbf{q})$ can be any arbitrary distance measure. The standard option is the total variational distance

$$d_1(\mathbf{p}, \mathbf{q}) = \frac{1}{2}\sum_i |p_i - q_i|.$$

With this choice of the distance, we have

Theorem 10. *For any w_1, w_2 satisfying $v \le w_2 \le w_1 \le w^*$, we have*

$$d_1(\mathcal{P}^c_{\mathtt{U}}(w_2), \mathcal{P}_{\mathtt{U}}(w_1)) \ge \frac{|w_1 - w_2|}{|\overline{m} - \underline{m}|}.$$

Observe that the quantities $\underline{m}$ and $\overline{m}$ can be simply computed from $\mathtt{U}$. The idea of the proof of Theorem 10 is standard (e.g., see (Gossner and Vieille, 2002, eq. (4)) for a similar derivation), but is included in Section 6.2 for completeness.

In this paper, we propose use of the Rényi divergence of order two between $\mathbf{p}$ and $\mathbf{q}$ to quantify the distance between two distributions:[4]

$$d_2(\mathbf{p}, \mathbf{q}) \triangleq \log(\sum_i \frac{p_i^2}{q_i}) = \log(1 + \chi^2(\mathbf{p}, \mathbf{q})).$$

Our first result gives the following bound:

Theorem 11. *For any w_1, w_2 satisfying $v \leq w_2 \leq w_1 \leq w^*$, we have*

$$d_2(\mathcal{P}^c_{\mathtt{U}}(w_2), \mathcal{P}_{\mathtt{U}}(w_1)) \geq \log\left(1 + \frac{(w_1 - w_2)^2}{(w_1 - \underline{m})(\overline{m} - w_1)}\right).$$

The proof can be found in Section 6.2. The above result is derived by a probabilistic approach, which we believe is novel in the context of linear programming.

Our second result is less crucial, but still useful. It gives a compact formula for the supporting hyperplanes of $\mathcal{P}_{\mathtt{U}}(w)$ in terms of the Nash equilibrium of a game. We need a definition:

Definition 12. *Let Nash($\mathtt{U}$) denote the value of a two-player zero-sum game with payoff table $\mathtt{U} = [u_{ij}]$; thus Nash($\mathtt{U}$) is the maximum value that Alice can guarantee using arbitrary mixed strategies. It is known that in a zero-sum game, while the game might have multiple Nash equilibriums, the value of Alice in all of the equilibriums is the same (Narahari, 2014, p.145).*

Given values $a_1, a_2, \cdots, a_n \in \mathbb{R}$, let us construct a new table whose (i, j) entry is $\tilde{u}_{ij} = u_{ij} + a_i$. In other words, we add a_i to the entries in the i-th row of $\mathtt{U}$. The new table can be expressed as $\tilde{\mathtt{U}} = \mathtt{U} + \mathbf{a} \cdot \mathbf{1}^T$ where $\mathbf{a}$ is a column vector whose entries are $a_1, a_2, \cdots, a_n$ and $\mathbf{1}^T$ is a row vector of all ones. Observe that the table $\tilde{\mathtt{U}}$ can be intuitively

[4] In this definition, we set p_i^2/q_i to be zero if $p_i = q_i = 0$, and infinity if $p_i > 0$ while $q_i = 0$. We have that $d_2(\mathbf{p}, \mathbf{q}) \geq 0$, and $d_2(\mathbf{p}, \mathbf{q}) = 0$ if and only if $\mathbf{p} = \mathbf{q}$.

understood as giving an additional incentive a_i to Alice for playing her i-th action (it is actually a disincentive or "tax" if $a_i < 0$).

Theorem 13. *The set $\mathcal{P}_{\mathtt{U}}(w)$ can be characterized as follows:*

$$\mathcal{P}_{\mathtt{U}}(w) = \left\{ \mathbf{p} \in \Delta(\mathcal{A}) \Big| \sum_i a_i p_i \leq \mathit{Nash}(\mathtt{U} + \mathbf{a}\mathbf{1}^T) - w, \quad \forall \mathbf{a} \right\}.$$

Remark 14. *Note that $\max_{\mathbf{p} \in \mathcal{P}_{\mathtt{U}}(w)} (\sum_i a_i p_i)$ is simply a linear program. Equivalence of linear program and Nash equilibriums are known in the literature Dantzig (1951); Adler (2013). However, our construction of the game $\tilde{\mathtt{U}}$ based on incentive or tax is new to best of our knowledge.*

In Section 6.3, we give the proof of Theorem 13 as well as a geometric picture of the Nash equilibrium strategies of Alice.

5.2. On the min-entropy function

We begin with a property of the min-entropy function. To state the property, we need the following definition:

Definition 15. *Consider two games with payoff matrixes $\mathtt{U}_1$ and $\mathtt{U}_2$. Let $\mathtt{U}_3 = \mathtt{U}_1 \oplus \mathtt{U}_2$ be the direct-sum of $\mathtt{U}_1$ and $\mathtt{U}_2$. Then $\mathtt{U}_3$ defines a new game in which players simultaneously play one instance of $\mathtt{U}_1$ and one instance of $\mathtt{U}_2$ and the resulting payoff is the sum of payoffs from $\mathtt{U}_1$ and $\mathtt{U}_2$.*

Theorem 16. *Then, we have $F_{\mathtt{U}\oplus\mathtt{U}}(w) = F_{\mathtt{U}}(w/2)$. Similarly, for every natural number k, $F_{\oplus^k\mathtt{U}}(w) = F_{\mathtt{U}}(w/k)$ where $\oplus^k\mathtt{U}$ is k times direct sum of $\mathtt{U}$.*

Proof of Theorem 16 is provided in Section 6.4.

An application of the above theorem is that given an expression $G_{\mathtt{U}}(w)$ that bounds $F_{\mathtt{U}}(w)$ from below for all w and $\mathtt{U}$, we can conclude that $G_{\oplus^k\mathtt{U}}(kw) \leq F_{\oplus^k\mathtt{U}}(kw) = F_{\mathtt{U}}(w)$. Thus, $\max_k G_{\oplus^k\mathtt{U}}(kw)$ is also a (potentially better) lower bound to $F_{\mathtt{U}}(w)$. As a result, we expect that a "good" lower (or upper bound) on $F_{\mathtt{U}}(w)$ should have the correct scaling behavior as we simultaneously play more and more copies of the game.

5.2.1. Lower and upper bounds on the min-entropy function

The min-entropy function $F_{\mathtt{U}}(w)$ in (9) is the minimum of a concave function on a polytope $\mathcal{P}_{\mathtt{U}}(w)$. This minimum occurs at a vertex of $\mathcal{P}_{\mathtt{U}}(w)$. This leads to a search in the exponentially large set of vertices of the polytope $\mathcal{P}_{\mathtt{U}}(w)$, which is computationally hard. We desire to find bounds on $F_{\mathtt{U}}(w)$ that are either explicit, or else can be computed in polynomial time. Observe that

$$F_{\mathtt{U}}(w) = \min_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} H(\mathbf{p}) = \log(|\mathcal{A}|) - \max_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} D(\mathbf{p}\|\mathbf{p}^{\mathtt{u}}) \tag{13}$$

where $\mathbf{p}^{\mathtt{u}}$ is the uniform distribution over $\mathcal{A}$, and $D(\mathbf{p}\|\mathbf{q}) = \sum_i p_i \log(p_i/q_i)$ is the Kullback–Leibler (KL) divergence. Thus, we are interested in finding the vertex of $\mathcal{P}_{\mathtt{U}}(w)$ which has maximum distance from the uniform distribution (with respect to KL divergence).

Lower bound: To prove lower bounds for

$$F_{\mathtt{U}}(w) = \min_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} H(\mathbf{p}), \tag{14}$$

one idea is to replace the entropy function with a smaller function and compute the minimum over $\mathcal{P}_{\mathtt{U}}(w)$. Another idea is to also relax the set of distributions $\mathcal{P}_{\mathtt{U}}(w)$ and replace it with a potentially bigger set. We proceed with the first idea, and comment on the second idea afterwards. Using the fact that the Rényi entropy is decreasing in its order, we obtain that for any $\alpha > 1$

$$F_{\mathtt{U}}(w) = \min_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} H(\mathbf{p}) \geq \min_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} H_\alpha(\mathbf{p}), \tag{15}$$

where $H_\alpha(\mathbf{p})$ is the Rényi entropy of order α:

$$H_\alpha(\mathbf{p}) = \frac{1}{1-\alpha} \log_2 \left(\sum_i p_i^\alpha \right).$$

The case of $\alpha = 2$ is related to the Euclidean norm and results in an optimization problem similar to the one given in (13) for the Euclidean norm instead of the KL divergence, which is still not tractable. However, the case $\alpha = \infty$ relates to the maximum norm and results

in the following lower bound:

$$G_{\mathtt{U}}^{(1)}(w) = -\log_2 \left(\max_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} \max_{i\in\mathcal{A}} p_i \right).$$

For each i, the problem of finding the maximum of p_i over $\mathbf{p} \in \mathcal{P}_{\mathtt{U}}(w)$ is a linear program. From Theorem 13, we can find an upper bound on the value of this linear program, yielding

$$G_{\mathtt{U}}^{(1)}(w) \geq -\log_2 \max_i \left(\mathrm{Nash}(\mathtt{U} + \mathbf{e}_i \mathbf{1}^T) - w \right), \tag{16}$$

where $\mathbf{e}_i$ is a vector of length $|\mathcal{A}| = n$ whose i-th coordinate is one, and all its other coordinates are zero. The lower bound $G_{\mathtt{U}}^{(1)}(w)$, or its relaxed version in (16) can be found in polynomial time, even though they are not in explicit forms.

To obtain an explicit lower bound, observe that $\log_2(1/p_i) = d_2(\mathbf{e}_i, \mathbf{p})$. Note that the vector $\mathbf{e}_i$ is a probability vector associated to a *deterministic* random variable that chooses i with probability one. Take some $\epsilon > 0$. By definition, deterministic strategies cannot secure a payoff of more than v. Hence $\mathbf{e}_i \in \mathcal{P}_{\mathtt{U}}^c(v+\epsilon)$. Thus,

$$\begin{aligned}
G_{\mathtt{U}}^{(1)}(w) &= -\log_2 \left(\max_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} \max_{i\in\mathcal{A}} p_i \right) \\
&= \min_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} \min_{i\in\mathcal{A}} \log_2(\frac{1}{p_i}) \\
&= \min_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} \min_{i\in\mathcal{A}} d_2(\mathbf{e}_i, \mathbf{p}) \\
&\geq \min_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} \min_{\mathbf{q}\in P_{\mathtt{U}}^c(v+\epsilon)} d_2(\mathbf{q}, \mathbf{p}) \\
&\geq \log_2 \left(1 + \frac{(w-v-\epsilon)^2}{(w-\underline{m})(\overline{m}-w)} \right), \qquad (17)
\end{aligned}$$

where (17) follows from Theorem 11. Letting $\epsilon \to 0$, we obtain

$$F_{\mathtt{U}}(w) \geq G_{\mathtt{U}}^{(2)}(w) \triangleq \log_2 \left(1 + \frac{(w-v)^2}{(w-\underline{m})(\overline{m}-w)} \right) \quad \forall w : v \leq w \leq w^*.$$

With a similar argument and using the fact that $p_i = 1 - d_1(\mathbf{e}_i, \mathbf{p})$, we obtain the following lower bound utilizing Theorem 10:

$$F_{\mathtt{U}}(w) \geq G_{\mathtt{U}}^{(3)}(w) \triangleq -\log_2 \left(1 - \frac{w-v}{\overline{m}-\underline{m}} \right) \quad \forall w : v \leq w \leq w^*.$$

Observe that when $v = \underline{m}$, $G_{\mathtt{U}}^{(2)}(w)$ equals $G_{\mathtt{U}}^{(3)}(w)$. When $v \neq \underline{m}$, a simple calculation shows that $G_{\mathtt{U}}^{(2)}(w) \geq G_{\mathtt{U}}^{(3)}(w)$ if and only if $w \geq (\overline{m} + v)/2$.

Example 17. *Consider a game with payoff matrix:*

$$\mathtt{U} = \begin{bmatrix} -1 & 1 & 1 \\ 1 & .5 & 1 \\ 1 & 1 & .5 \end{bmatrix}$$

From payoff matrix we have $v = .5$, $\underline{m} = -1$, $\overline{m} = 1$ *and* $w^* = 0.778$. *Therefore for* $w \geq 0.75$, $G_{\mathtt{U}}^{(2)}(w)$ *gives a better lower bound than* $G_{\mathtt{U}}^{(3)}(w)$ *on* $F_{\mathtt{U}}(w)$.

Remark 18. *One can inspect that just like that the min-entropy function, the explicit lower bounds* $G_{\mathtt{U}}^{(2)}(w)$ *and* $G_{\mathtt{U}}^{(3)}(w)$ *satisfy*

$$G_{\mathtt{U}\oplus\mathtt{U}}^{(i)}(w) = G_{\mathtt{U}}^{(i)}(w/2), \quad i = 2, 3,$$

have the correct scaling behavior. Additionally, by replacing entropy with Rényi entropy in the proof of Theorem 16, one also obtains that

$$G_{\mathtt{U}\oplus\mathtt{U}}^{(1)}(w) = G_{\mathtt{U}}^{(1)}(w/2).$$

The function $G_{\mathtt{U}}^{(2)}(w)$ *has second derivative for all* $v \leq w \leq w^*$, *while the second derivative of the piecewise concave function* $F_{\mathtt{U}}(w)$ *is defined everywhere except for a finite number of kink points. The second derivative of the function* $G_{\mathtt{U}}^{(2)}(w)$ *may be positive or negative, while* $F_{\mathtt{U}}(w)$ *is piecewise concave. On the other hand, the function* $G_{\mathtt{U}}^{(1)}(w)$ *is piecewise convex. The reason is that if* $\max_{i\in\mathcal{A}} \max_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} p_i$ *is attained by* i^* *and a particular vertex of* $\mathcal{P}_{\mathtt{U}}(w)$ *for* $w \in [w_1, w_2]$, *in this interval* $\max_{i\in\mathcal{A}} \max_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} p_i$ *varies linearly in* w. *Then, convexity of* $-\log_2(\cdot)$ *results in convexity of* $G_{\mathtt{U}}^{(1)}(w)$ *in the interval* $[w_1, w_2]$.

Remark 19. *Inequality* (17) *shows that*

$$\max_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} p_i \leq \left(1 + \frac{(w-v)^2}{(w-\underline{m})(\overline{m}-w)}\right)^{-1} \quad \forall w : v \leq w \leq w^* \text{and } \forall i = 1, \ldots, n \tag{18}$$

which gives an upper bound for the linear programming of $\max_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} p_i$. *Note that* $\mathcal{P}_{\mathtt{U}}(w)$ *is a very generic polytope, parameterized by a variable* w. *By a change of variables (scaling and shifting), one can convert* $\max_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w)} p_i$ *to a wide class of linear programs (with no immediate connection to the probability simplex), and then use the bound given in* (18).

Let us now turn to the second idea to prove a lower bound for $F_{\mathtt{U}}(w)$ in (14), namely replacing the set of distributions $\mathcal{P}_{\mathtt{U}}(w)$ with a potentially bigger set. As mentioned earlier, minimization of the entropy over the set

$$\mathcal{P}_{\mathtt{U}}(w) = \{\mathbf{p} \in \Delta(\mathcal{A}) : \mathbf{p}^T\mathtt{U} \geq w\mathbf{1}^T\}$$

can be difficult. However, it could be possible to solve it (or find good lower bounds for it) for special choices of the matrix $\mathtt{U}$.[5] We show how a result for an special case of $\mathtt{U}$ can be utilized to find a bound (computable in polynomial time) for an arbitrary $\mathtt{U}$. In particular, let us assume that one has a way to minimize entropy over the set

$$\mathcal{Q}(\mathbf{r}) \triangleq \{\mathbf{p} \in \Delta(\mathcal{A}) : \mathbf{p}^T\mathtt{U}^* \geq \mathbf{r}^T\}$$

for some given matrix $\mathtt{U}^*$, and any arbitrary column vector $\mathbf{r}$. Then, we are interested in a value for $\mathbf{r}$ such that

$$\mathcal{P}_{\mathtt{U}}(w) \subseteq \mathcal{Q}(\mathbf{r}). \tag{19}$$

Then, we can relax minimization of the entropy over the set $\mathcal{P}_{\mathtt{U}}(w)$ by instead, computing its minimum over the set $\mathcal{Q}(\mathbf{r})$. Note that an appropriate $\mathbf{r}$ in (19) can be found by solving a number of linear programs: the product $\mathbf{p}^T\mathtt{U}^*$ consists of a number of linear equations on coordinates of $\mathbf{p}$, and the minimum of each linear equation over the set $\mathcal{P}_{\mathtt{U}}(w)$ is a linear program (see also Theorem 13).

Upper bound: It is clear that any $F_{\mathtt{U}}(w) \leq H(\mathbf{p})$ for any arbitrary choice of $\mathbf{p} \in \mathcal{P}_{\mathtt{U}}(w)$. The following theorem gives a number of upper bounds, each of which are obtained by identifying $\mathbf{p} \in \mathcal{P}_{\mathtt{U}}(w)$ in different ways.

[5]For instance, if each row of $\mathtt{U}$ has only one non-zero element, the set of constraints will be on the individual coordinates of the vector $\mathbf{p}$ and minimizing entropy for such constraints is tractable.

Theorem 20. *Consider a game with payoff matrix* $\mathtt{U}$ *and parameters* v, $\underline{m}$, $\overline{m}$ *and* w^* *defined in Definition 9. Let* h^* *be the entropy of some Nash strategy of player A and define*

$$Q_{\mathtt{U}}^{(1)}(w) = \min\left\{\mathsf{h}^*, \frac{w-v}{w^*-v}\mathsf{h}^* + h\left(\frac{w-v}{w^*-v}\right)\right\}$$

$$Q_{\mathtt{U}}^{(2)}(w) = \min_{i\in\mathcal{A}} H(\mathbf{p}^*_{max,i}), \quad \mathbf{p}^*_{max,i} \in \arg\max_{\mathbf{p}\in\mathcal{P}(w)} \quad p_i$$

$$Q_{\mathtt{U}}^{(3)}(w) = \min_{j\in\mathcal{X}_{\mathcal{B}}} \quad \max_{\mathbf{p}\in\mathcal{P}_{\mathtt{U}}(w):\mathbf{p}^T\mathtt{U}\mathbf{e}_j=w} H(\mathbf{p})$$

where $h(\alpha) = -\alpha\log(\alpha) - (1-\alpha)\log(1-\alpha)$ *and* $Q_{\mathtt{U}}^{(2)}(w)$ *can be defined with any choice of* $\mathbf{p}^*_{max,i}$ *from the argmax set (if there are multiple possible choices). Therefore,*

$$F_{\mathtt{U}}(w) \le Q_{\mathtt{U}}^{(r)}(w), \quad r = 1, 2, 3.$$

Proof of Theorem 20 can be found in Section 6.5.

Remark 21. *The second derivative of* $Q_{\mathtt{U}}^{(1)}(w)$ *is negative for* $v \le w \le w^*$. *Thus* $Q_{\mathtt{U}}^{(1)}(w)$ *is a concave function of* w *and one can readily inspect that:*

$$Q_{\mathtt{U}\oplus\mathtt{U}}^{(1)}(w) = Q_{\mathtt{U}}^{(1)}(\frac{w}{2}).$$

As $\arg\max_{\mathbf{p}\in\mathcal{P}(w)} p_i$ *may contain multiple elements,* $Q_{\mathtt{U}}^{(2)}(w)$ *is not a well defined function of* w. *The function* $Q_{\mathtt{U}}^{(3)}(w)$ *is not necessarily scalable for game* $\mathtt{U}\oplus\mathtt{U}$.

Example 22. *Consider a game with payoff matrix:*

$$\mathtt{U} = \begin{bmatrix} 3 & 1 & 0 & -2 & 0 & -2 \\ 1 & 3 & -2 & 0 & -2 & 0 \\ 0 & -2 & 3 & 1 & 0 & -2 \\ -2 & 0 & 1 & 3 & -2 & 0 \\ 0 & -2 & 0 & -2 & 3 & 1 \\ -2 & 0 & -2 & 0 & 1 & 3 \end{bmatrix}$$

Figure 3 illustrates the behavior of the bounds for this example. In this example $v = \underline{m}$; *hence* $G_{\mathtt{U}}^{(3)}(w)$ *coincides with* $G_{\mathtt{U}}^{(2)}(w)$ *and is not depicted in the figure.*

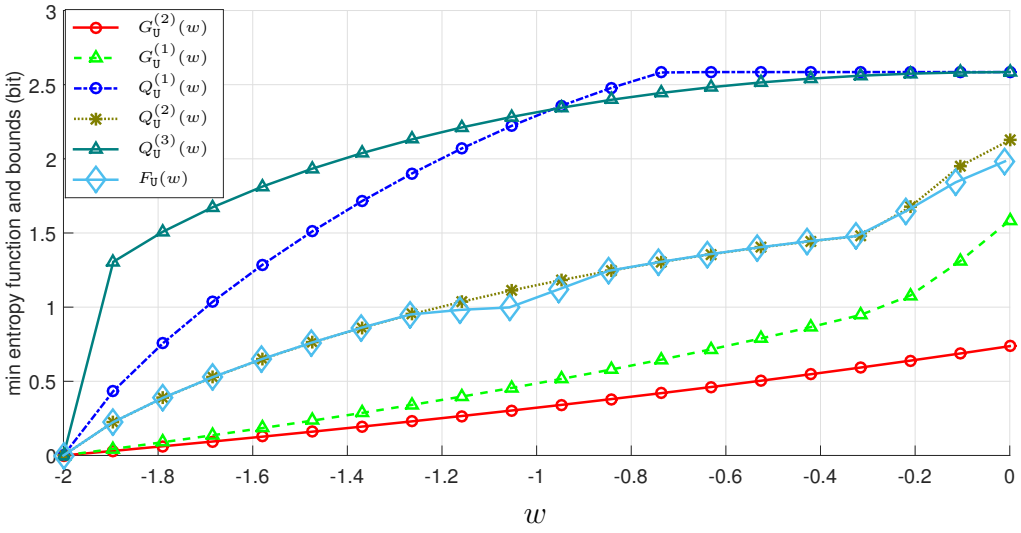

Figure 3: Illustration of the bounds on min-entropy function for the game defined in example 22. The horizontal line depicts w and vertical line depicts the value of bounds.

6. Proofs

6.1. Proof of Proposition 8

To prove Proposition 8 we make use of Lemma 23 which is a direct consequence of (Yassaee et al., 2014, Theorem 1) and Lemma 2.

Lemma 23. *Consider the correlated random sequences X^L and Y^L drawn i.i.d. from respective spaces $\mathcal{X}$ and $\mathcal{Y}$ by joint probability distribution $p_{X,Y}$. Let $Q^{(L)}$ be a random variable independent of Y^L and uniformly distributed on $\{1, 2, \ldots, 2^{RL}\}$ where $R < H(X|Y)$ is a real number, then there exists a mapping $\mathsf{B} : \mathcal{X}^L \to \{1, 2, \ldots, 2^{RL}\}$ such that*

$$\lim_{L\to\infty} \|p_{\mathsf{B}(X^L)Y^L} - p_{Q^{(L)}Y^L}\|_{TV} = 0.$$

Choose a real number R such that $\gamma H(p_A^{(1)}) + (1-\gamma)H(p_A^{(2)}) < R < H(X|Y)$ and LR is a natural number (for a sufficiently large L).

We define random variables $Q^{(L)}$ and $\hat{A}^L$ that are mutually independent of each other, and of Y^L with the following marginal distributions: let $Q^{(L)}$ be a uniformly distributed

random variable on $\{1, 2, \ldots, 2^{RL}\}$, and $\hat{A}^L$ be distributed as follows:

$$p_{\hat{A}^L}(a^L) = \prod_{t=1}^{\lceil \gamma L \rceil} p_A^{(1)}(a_t) \prod_{t=\lceil \gamma L) \rceil + 1}^{L} p_A^{(2)}(a_t)$$

where $\lceil a \rceil$ is the smallest integer greater than or equal to a. Note that $R < H(X|Y)$; thus according to Lemma 23 there exists a mapping $\mathtt{B} : \mathcal{X}^L \to \{1, 2, \ldots, 2^{RL}\}$ satisfying

$$\lim_{L \to \infty} \| p_{\mathtt{B}(X^L)Y^L} - p_{Q^{(L)}Y^L} \|_{TV} = 0. \tag{20}$$

On the other hand we have

$$\begin{aligned} \text{p-}\limsup_{L \to \infty} \frac{1}{L} \log \frac{1}{p_{\hat{A}^L}(\hat{A}^L)} &= \gamma H(p_A^{(1)}) + (1-\gamma) H(p_A^{(2)}) \\ &< R = \text{p-}\liminf_{L \to \infty} \frac{1}{L} \log \frac{1}{p_{Q^L}(Q^L)}, \end{aligned}$$

therefore according to Lemma 2 there exists a mapping $\varphi : \{1, 2, \ldots, 2^{RL}\} \to \mathcal{A}^L$ such that

$$\lim_{L \to \infty} \| p_{\varphi\left(Q^{(L)}\right)} - p_{\hat{A}^L} \|_{TV} = 0.$$

Considering the fact that $\varphi(Q^{(L)})$ and $\hat{A}^L$ both are independent of Y^L, the above equation results in

$$\lim_{L \to \infty} \| p_{\varphi\left(Q^{(L)}\right)Y^L} - p_{\hat{A}^L Y^L} \|_{TV} = 0. \tag{21}$$

Furthermore note that

$$\begin{aligned} \| p_{\varphi(\mathtt{B}(X^L))Y^L} - p_{\hat{A}^L Y^L} \|_{TV} &\leq \| p_{\varphi(\mathtt{B}(X^L))Y^L} - p_{\varphi\left(Q^{(L)}\right)Y^L} \|_{TV} + \| p_{\varphi\left(Q^{(L)}\right)Y^L} - p_{\hat{A}^L Y^L} \|_{TV} \\ &\leq \| p_{\mathtt{B}(X^L)Y^L} - p_{Q^{(L)}Y^L} \|_{TV} + \| p_{\varphi\left(Q^{(L)}\right)Y^L} - p_{\hat{A}^L Y^L} \|_{TV} \end{aligned} \tag{22}$$

where the first inequality follows from the triangle inequality for the total variation distance, and the second inequality holds because

$$\| p_{f(Y)} - p_{f(Z)} \|_{TV} \leq \| p_Y - p_Z \|_{TV}$$

for any arbitrary deterministic function $f(\cdot)$ on the sample space of r.v.s Y and Z. Let $A^L = \psi(X^L) = \varphi\left(\mathtt{B}\left(X^L\right)\right)$. Then by combining equations (20), (21) and (22) we have

$$\lim_{L \to \infty} \| p_{A^L Y^L} - p_{\hat{A}^L Y^L} \|_{TV} = 0. \tag{23}$$

Thus $\psi(\cdot) = \varphi(\mathtt{B}(\cdot))$ is the desired mapping.

6.2. Proof of Theorems 10 and 11

Take arbitrary $\mathbf{p} \in \mathcal{P}^c_{\mathtt{U}}(w_2)$ and $\mathbf{q} \in \mathcal{P}_{\mathtt{U}}(w_1)$. Since $\mathbf{p} \in \mathcal{P}^c_{\mathtt{U}}(w_2)$, when Alice plays according to $\mathbf{p}$, Bob has a (pure) strategy $j \in \mathcal{B}$ that reduces Alice's expected payoff to a number less than w_2. We can define a random variable $\tilde{W}$ that represents the payoff of Alice when Alice plays $\mathbf{p}$ and Bob plays j. Then, the alphabet set of $\tilde{W}$ is $\{u_{1j}, u_{2j}, \cdots, u_{nj}\}$, and its probability distribution over this set is $\mathbf{p}$. We must have $\mathbb{E}[\tilde{W}] < w_2$.

Let us assume that Alice adopts $\mathbf{q}$, but Bob keeps playing the same $j \in \mathcal{B}$. Let W be the random variable, describing the payoff of Alice when she plays according to $\mathbf{q}$. Because $\mathbf{q} \in \mathcal{P}_{\mathtt{U}}(w_1)$, $\mathbb{E}[W] \geq w_1$. We can think of W as taking u_{ij} with probability p_i.

Theorem 10 follows from the following chain of inequalities:

$$\begin{aligned}
|\mathbb{E}[W] - \mathbb{E}[\tilde{W}]| &= \left| \sum_i (p_i - q_i) u_{ij} \right| \\
&= \left| \sum_i (p_i - q_i) \left(u_{ij} - \frac{\overline{m} + \underline{m}}{2} \right) \right| \\
&\leq \Big(\sum_i |(p_i - q_i)| \Big) \max_i \left| u_{ij} - \frac{\overline{m} + \underline{m}}{2} \right| \\
&\leq (2 d_1(\mathbf{p}, \mathbf{q})) \left(\frac{|\overline{m} - \underline{m}|}{2} \right). \qquad (24)
\end{aligned}$$

To prove Theorem 11, the key step is to relate the Rényi divergence to *variance*. Remember that $d_2(\mathbf{p}, \mathbf{q}) = \log(1 + \chi^2(\mathbf{p}, \mathbf{q}))$.

Lemma 24. *Take an arbitrary sequence of real numbers $(x, y, \cdots, x_n)$ and two probability distributions $\mathbf{p} = (p_1, p_2, \cdots, p_n)$ and $\mathbf{q} = (q_1, q_2, \cdots, q_n)$. Let W be a random variable that takes value x_i with probability q_i, and $\tilde{W}$ be a random variable that takes value x_i with probability p_i. Then, we have*

$$\chi^2(\mathbf{p}, \mathbf{q}) \geq \frac{(\mathbb{E}[W] - \mathbb{E}[\tilde{W}])^2}{\mathsf{Var}[W]}. \qquad (25)$$

Furthermore, the above inequality becomes an equality if we set $x_i = (p_i - q_i)/q_i$.

Observe that the left hand side of (25) depends only on the probability values p_i and q_i, while the right hand side depends not only on the probabilities, but also the values that W and $\tilde{W}$ take. The proof of this lemma is given in Appendix A. Using this lemma and the fact that $\mathbb{E}[\tilde{W}] < w_2 < w_1 \leq \mathbb{E}[W]$, we can conclude

$$\chi^2(\mathbf{p}, \mathbf{q}) \geq \frac{(\mathbb{E}[W] - w_2)^2}{\mathsf{Var}[W]}.$$

Observe that $\underline{m} \leq W \leq \overline{m}$ holds with probability one. The proof is finished by the following lemma.

Lemma 25. *For any $w_1 > w_2$, we have*

$$\frac{(\mathbb{E}[W] - w_2)^2}{\mathsf{Var}[W]} \geq \frac{(w_1 - w_2)^2}{(w_1 - \underline{m})(\overline{m} - w_1)},$$

provided that $\mathbb{E}[W] \geq w_1$, and $\underline{m} \leq W \leq \overline{m}$.

The proof of this lemma is given in Appendix B.

6.3. Proof of Theorem 13

In this section we provide additional details and build a geometric picture. This picture implies Theorem 13, but also gives a geometric interpretation of Nash strategies.

Let

$$L(a_1, a_2, \cdots, a_n) = \mathrm{Nash}(\mathtt{U} + \mathbf{a}\mathbf{1}^T), \qquad \forall \mathbf{a} \in \mathbb{R}^n.$$

Observe that the table $\tilde{\mathtt{U}}$ can be intuitively understood as giving an additional incentive a_i to Alice for playing her i-th action (it is actually a disincentive if $a_i < 0$). Also, since

$$L(a_1 + c, a_2 + c, \cdots, a_n + c) = L(a_1, a_2, \cdots, a_n) + c,$$

we only need to understand L when the sum of the incentives a_i is zero.

We need the following definition:

Definition 26. *Let*

$$K(p_1, p_2, \cdots, p_n) = \min_j \sum_i p_i u_{ij}$$

be the payoff that Alice can guarantee with playing distribution $(p_1, \cdots, p_n)$ *with table* $\mathtt{U}$. *We extend the definition of* $K(\cdot)$ *to arbitrary* $(p_1, p_2, \cdots, p_n) \in \mathbb{R}^n$ *by setting*

$$K(p_1, p_2, \cdots, p_n) = -\infty,$$

when the tuple $(p_1, p_2, \cdots, p_n)$ *does not form a valid probability distribution,* i.e., *when any of the* p_i*'s becomes negative, or* $\sum_i p_i \neq 1$.

Note that

$$\mathcal{P}_{\mathtt{U}}(w) = \left\{ \mathbf{p} \middle| K(\mathbf{p}) \geq w \right\}.$$

A full geometric picture of $\mathcal{P}_{\mathtt{U}}(w)$ as well as Nash strategies are provided in the following theorem:

Theorem 27. *We have*

1. *The function* $L(a_1, a_2, \cdots, a_n)$ *is the convex conjugate dual of* $K(p_1, p_2, \cdots, p_n)$ *in the following sense:*

$$L(\mathbf{a}) = \max_{\mathbf{p} \in \mathbb{R}^n} \left[K(\mathbf{p}) + \sum_{i=1}^{n} p_i a_i \right], \qquad \forall \mathbf{a} \in \mathbb{R}^n. \tag{26}$$

   The function $L(a_1, a_2, \cdots, a_n)$ *is jointly convex in* $(a_1, a_2, \cdots, a_n)$, *while* $K(p_1, p_2, \cdots, p_n)$ *is jointly concave in* $(p_1, p_2, \cdots, p_n)$.

   Furthermore, the supporting hyperplanes to the convex curve $\mathbf{a} \mapsto L(\mathbf{a})$ *characterize Alice's Nash strategies as follows: for any arbitrary* $\mathbf{a}$, $\mathbf{p}$ *is a Nash strategy of Alice for table* $\tilde{\mathtt{U}} = \mathtt{U} + \mathbf{a}\mathbf{1}^T$ *if and only if* $\mathbf{p}$ *is a subgradient of the function* L *at* $\mathbf{a}$. *In other words, take some arbitrary vector* $\mathbf{a}$. *Then,*

$$L(b_1, b_2, \cdots, b_n) \geq L(a_1, a_2, \cdots, a_n) + \sum_i (b_i - a_i) p_i, \qquad \forall \mathbf{b} \in \mathbb{R}^n, \tag{27}$$

   if and only if $\mathbf{p}$ *is a Nash strategy of Alice for the payoff table* $\tilde{\mathtt{U}} = \mathtt{U} + \mathbf{a}\mathbf{1}^T$.

2. *Given a probability vector* $\mathbf{p}$, *we have*

$$L(a_1, a_2, \cdots, a_n) \geq w + \sum_i (a_i - b_i) p_i, \qquad \forall \mathbf{a},$$

if and only if $\mathbf{p}$ *guarantees a payoff of at least* w *for game* $\mathtt{U} + \mathbf{b}\mathbf{1}^T$*. In particular, setting* $b_i = 0$*,*

$$L(a_1, a_2, \cdots, a_n) \geq w + \sum_i a_i p_i, \qquad \forall \mathbf{a}, \tag{28}$$

if and only if $\mathbf{p}$ *guarantees a payoff of at least* w *for game* $\mathtt{U}$*,* i.e., $\mathbf{p} \in \mathcal{P}_{\mathtt{U}}(w)$*. Thus, having a payoff* w*, we look for hyperplanes of the form* $w + \sum_i p_i a_i$ *that pass through* w *at* $(a_1, a_2, \cdots, a_n) = (0, 0, \cdots, 0)$*, and lie below the curve of* L*.*

Observe that the second part of Theorem 27 is equivalent with Theorem 13.

Proof of Theorem 27. We begin with the first part of the theorem. Using the max-min formulation for the Nash value of a game $\tilde{\mathtt{U}}$, we have

$$\begin{aligned}
L(a_1, a_2, \cdots, a_n) &= \max_{p_i \geq 0, \sum_i p_i = 1} \min_j \sum_i p_i (u_{ij} + a_i) \\
&= \max_{p_i \geq 0, \sum_i p_i = 1} \left[\left(\min_j \sum_i p_i u_{ij} \right) + \sum_i p_i a_i \right] \\
&= \max_{p_i \geq 0, \sum_i p_i = 1} \left[K(p_1, p_2, \cdots, p_n) + \sum_i p_i a_i \right]. \\
&= \max_{p_i \in \mathbb{R}} \left[K(p_1, p_2, \cdots, p_n) + \sum_i p_i a_i \right]. \qquad (29)
\end{aligned}$$

where (29) follows from the fact that $K(\mathbf{p})$ is minus infinity when $\mathbf{p}$ is not a probability distribution. This shows the duality of $L(\cdot)$ and $K(\cdot)$. Next, note that $K(\cdot)$ is a minimum of linear functions; hence it is a concave function. Convexity of $L(\cdot)$ can be directly seen from (27) which implies that at least one supporting hyperplane to its curve exists at any given point (since at least one Nash strategy exists for any arbitrary game). Thus, it remains to prove (27).

Without loss of generality, it suffices to prove (27) for $\mathbf{b} = \mathbf{0}$ ($b_i = 0$), and get the result for arbitrary $\mathbf{b}$ by changing variable $\mathtt{U} \to \mathtt{U} + \mathbf{b}\mathbf{1}^T$. The inequality

$$L(a_1, a_2, \cdots, a_n) \geq L(0, 0, \cdots, 0) + \sum_i a_i p_i, \qquad \forall \mathbf{a}, \tag{30}$$

can be also expressed as

$$\min_{\mathbf{a}} \left(L(a_1, a_2, \cdots, a_n) - \sum_i a_i p_i \right) \geq L(0, 0, \cdots, 0).$$

From the duality relation (29) and utilizing the Fenchel's duality theorem, the left hand side is $K(p_1, p_2, \cdots, p_n)$. Thus, the expression is equivalent with (30) can be written as

$$K(p_1, p_2, \cdots, p_n) \geq L(0, 0, \cdots, 0) = \text{Nash}(\mathtt{U}),$$

which is equivalent with $\mathbf{p}$ being a Nash strategy.

To proof for the second part of the theorem is similar. As before b_i can be set to zero. Then, we can express (28) as

$$\min_{\mathbf{a}} \left(L(a_1, a_2, \cdots, a_n) - \sum_i a_i p_i \right) \geq w.$$

From the duality relation (29) and the Fenchel's duality theorem, this is equivalent with $K(p_1, p_2, \cdots, p_n) \geq w$, or $\mathbf{p} \in \mathcal{P}_{\mathtt{U}}(w)$.

□

6.4. Proof of Theorem 16

Let X_1 and X_2 be two r.v.s representing the actions of player A in game $\mathtt{U} \oplus \mathtt{U}$ and $p(x_1)$ be an arbitrary distribution that secures payoff $w/2$ in game $\mathtt{U}$ then linearity of expectation results that $p(x_1, x_2) = \mathbb{1}(x_1 = x_2)p(x_1)$ secures payoff w in game $\mathtt{U} \oplus \mathtt{U}$, where $\mathbb{1}(\cdot)$ is indicator function. In this case $H(X_1, X_2) = H(X_1)$ thus the function $F_{\mathtt{U}\oplus\mathtt{U}}(w)$ is bounded above as follows:

$$F_{\mathtt{U}\oplus\mathtt{U}}(w) \leq F_{\mathtt{U}}(\frac{w}{2}). \tag{31}$$

On the other hand, let $p^*(x_1, x_2) \in \mathcal{P}_{\mathtt{U}\oplus\mathtt{U}}(w)$ be a distribution with minimum entropy that secures a payoff w in game $\mathtt{U} \oplus \mathtt{U}$, *i.e.,* $H(p^*(x_1, x_2)) = F_{\mathtt{U}\oplus\mathtt{U}}(w)$. Note that such a

distribution $p^*(x_1, x_2)$ exist because $\mathcal{P}_{\mathtt{U}\oplus\mathtt{U}}(w)$ is compact and closed. We have

$$w \leq \min_{j_1,j_2\in\mathcal{B}} \mathbb{E}_{p^*(x_1,x_2)}[u_{X_1,j_1} + u_{X_2,j_2}] \tag{32}$$

$$= \min_{j_1,j_2\in\mathcal{B}} \left(\mathbb{E}_{p^*(x_1,x_2)}[u_{X_1,j_1}] + \mathbb{E}_{p^*(x_1,x_2)}[u_{X_2,j_2}]\right) \tag{33}$$

$$= \min_{j_1\in\mathcal{B}} \mathbb{E}_{p^*(x_1)}[u_{X_1,j_1}] + \min_{j_2\in\mathcal{B}} \mathbb{E}_{p^*(x_2)}[u_{X_2,j_2}] \tag{34}$$

As a result, there exists $k \in \{1, 2\}$ such that

$$\min_{j_k\in\mathcal{B}} \mathbb{E}_{p^*(x_k)}[u_{X_k,j_k}] \geq \frac{w}{2}. \tag{35}$$

Thus, $p^*(x_k)$ secures payoff $w/2$ in game $\mathtt{U}$. Hence,

$$F_{\mathtt{U}}(w) = H(p^*(x_1, x_2)) \overset{(a)}{\geq} H(p^*(x_k)) \geq F_{\mathtt{U}}(\frac{w}{2}), \tag{36}$$

where (a) follows from properties of entropy.

Equations (31) and (36) conclude $F_{\mathtt{U}\oplus\mathtt{U}}(w) = F_{\mathtt{U}}(w/2)$. The above line of proof can be extended for every natural number k to prove $F_{\oplus^k\mathtt{U}}(w) = F_{\mathtt{U}}(w/k)$.

6.5. Proof of Theorem 20

Let $\mathbf{p}^*$ be a Nash strategy for player A where $H(\mathbf{p}^*) = \mathsf{h}^*$, and $i^* \in \arg\max_{i\in\mathcal{A}} \min_{j\in\mathcal{B}} u_{ij}$ be the pure strategy that guarantees payoff v for player A. For $v \leq w \leq w^*$ define $\alpha = (w^* - w)/(w^* - v)$ and $\mathbf{p} = \alpha\mathbf{e}_{i^*} + (1 - \alpha)\mathbf{p}^*$ where $\mathbf{e}_{i^*}$ is a vertical vector of all zero elements except for its i^*th element which is 1. The pmf $\mathbf{p}$ guarantees payoff w since

$$\mathbf{p}^T U = \alpha\mathbf{e}_{i^*}^T U + (1 - \alpha)\mathbf{p}^{*T} U \geq \alpha v\mathbf{1} + (1 - \alpha)w^*\mathbf{1} = w\mathbf{1}$$

Using properties of entropy one can inspect that:

$$H(\mathbf{p}) \leq \alpha H(\mathbf{e}_{i^*}) + (1-\alpha)H(\mathbf{p}^*) + H(\alpha, 1-\alpha) = (1-\alpha)\mathsf{h}^* - \alpha\log_2(\alpha) - (1-\alpha)\log_2(1-\alpha).$$

Therefore $F_{\mathtt{U}}(w) \leq H(\mathbf{p}) \leq (1 - \alpha)\mathsf{h}^* - \alpha\log_2(\alpha) - (1 - \alpha)\log_2(1 - \alpha)$. Additionally, as h^* secures payoff $w^* \geq w$ then $F_{\mathtt{U}}(w) \leq \mathsf{h}^*$. These two facts imply $F_{\mathtt{U}}(w) \leq Q^{(1)}_{\mathtt{U}}(w)$ once we substitute the value of α.

From the definition of $\mathbf{p}^*_{\max,i}$ it follows that $\mathbf{p}^*_{\max,i} \in \mathcal{P}_{\mathtt{U}}(w)$. Therefore for every $i \in \mathcal{A}$, $F_{\mathtt{U}}(w) \leq H(\mathbf{p}^*_{\max,i})$ and $F_{\mathtt{U}}(w) \leq Q^{(2)}_{\mathtt{U}}(w)$ follows.

For an arbitrary $j \in \mathcal{B}$ if $\mathbf{p}' \in \{\mathbf{p} \in \mathcal{P}_{\mathtt{U}}(w) : \mathbf{p}^T \mathtt{U} \mathbf{e}_j = w\}$ then as $\mathbf{p}'$ also belongs to $\mathcal{P}_{\mathtt{U}}(w)$ we have $F_{\mathtt{U}}(w) \leq H(\mathbf{p}')$. Then, for every $j \in \mathcal{B}$:

$$F_{\mathtt{U}}(w) \leq \max_{\mathbf{p} \in \mathcal{P}_{\mathtt{U}}(w) : \mathbf{p}^T \mathtt{U} \mathbf{e}_j = w} H(\mathbf{p})$$

As the above inequality is correct for every $j \in \mathcal{B}$, $F_{\mathtt{U}}(w) \leq Q^{(3)}_{\mathtt{U}}(w)$.

Appendix A. Proof of Lemma 24

We would like to prove that

$$\chi^2(\mathbf{p}, \mathbf{q}) \geq \frac{(\mathbb{E}[W] - \mathbb{E}[\tilde{W}])^2}{\mathsf{Var}[W]}, \tag{A.1}$$

where $W \sim \mathbf{q}$ and $\tilde{W} \sim \mathbf{p}$ are defined on some arbitrary set $\mathcal{X}$. To the best of our knowledge, this inequality has not appeared in the game theory or information theory literature before. We had a direct proof of it, until we found out that it can be derived from a result of Broniatowski and Leorato (2006) published in a math journal in 2006. According to (Broniatowski and Leorato, 2006, Lemma 1), we have

$$\chi^2(\mathbf{p}, \mathbf{q}) \geq \mathbb{E}[W] - \mathbb{E}[\tilde{W}] - \frac{1}{4}\mathbb{E}[W^2]. \tag{A.2}$$

Since (A.2) holds for all W and $\tilde{W}$, let us apply the transformation $W \mapsto a(W - \mathbb{E}[W])$, and $\tilde{W} \mapsto a(\tilde{W} - \mathbb{E}[W])$. This shows that for any real a, we have

$$\chi^2(\mathbf{p}, \mathbf{q}) \geq a\left(\mathbb{E}[W] - \mathbb{E}[\tilde{W}]\right) - \frac{a^2}{4}\mathsf{Var}[W]. \tag{A.3}$$

The right hand side is a quadratic expression in terms of a, and reaches its maximum at

$$a = \frac{2\left(\mathbb{E}[W] - \mathbb{E}[\tilde{W}]\right)}{\mathsf{Var}[W]}.$$

Substituting this value of a into (A.3), we get the desired result. The second part of the theorem that (A.1) becomes an equality for $x_i = (p_i - q_i)/q_i$ can be verified directly.

Appendix B. Proof of Lemma 25

Let us minimize the expression $(\mathbb{E}[X] - w_2)^2/\mathsf{Var}[X]$ over all r.v.s X that satisfy $\underline{m} \leq X \leq \overline{m}$ and $\mathbb{E}[X] \geq w_1$:

$$\min_{\substack{\underline{m}\leq X\leq \overline{m}\\ \mathbb{E}[X]\geq w_1}} \frac{\left(\mathbb{E}[X]-w_2\right)^2}{\mathsf{Var}[X]} = \min_{w_1\leq \mu\leq \overline{m}} \left(\min_{\substack{\underline{m}\leq X\leq \overline{m}\\ \mathbb{E}[X]=\mu}} \frac{\left(\mathbb{E}[X]-w_2\right)^2}{\mathsf{Var}[X]} \right) = \min_{w_1\leq \mu\leq \overline{m}} \left(\frac{(\mu-w_2)^2}{\max\limits_{\substack{\underline{m}\leq X\leq \overline{m}\\ \mathbb{E}[X]=\mu}} \mathsf{Var}[X]} \right)$$

$$= \min_{w_1\leq \mu\leq \overline{m}} \left(\frac{(\mu-w_2)^2}{-\mu^2 + \max\limits_{\substack{\underline{m}\leq X\leq \overline{m}\\ \mathbb{E}[X]=\mu}} \mathbb{E}[X^2]} \right). \tag{B.1}$$

We claim that

$$\max_{\substack{\underline{m}\leq X\leq \overline{m}\\ \mathbb{E}[X]=\mu}} \mathbb{E}[X^2] = (\overline{m}-\mu)(\mu-\underline{m}) + \mu^2. \tag{B.2}$$

Observe that if X^* is the following binary random variable

$$\mathbb{P}[X^* = \underline{m}] = \frac{\overline{m}-\mu}{\overline{m}-\underline{m}}, \quad \mathbb{P}[X^* = \overline{m}] = \frac{\mu-\underline{m}}{\overline{m}-\underline{m}} \tag{B.3}$$

we have $\mathbb{E}[X^*] = \mu$ and $\mathbb{E}[(X^*)^2] = (\overline{m}-\mu)(\mu-\underline{m}) + \mu^2$. As a result,

$$\max_{\substack{\underline{m}\leq X\leq \overline{m}\\ \mathbb{E}[X]=\mu}} \mathbb{E}[X^2] \geq (\overline{m}-\mu)(\mu-\underline{m}) + \mu^2. \tag{B.4}$$

On the other hand, the function $f(x) = x^2$ is convex and lies below the line that connects the two points $(\underline{m}, \underline{m}^2)$ and $(\overline{m}, \overline{m}^2)$ for any $x \in [\underline{m}, \overline{m}]$, *i.e.,*

$$x^2 \leq \underline{m}^2 + (x-\underline{m})(\overline{m}+\underline{m}), \quad \forall x \in [\underline{m}, \underline{m}^2].$$

Thus,

$$\mathbb{E}[X^2] \leq \underline{m}^2 + (\mathbb{E}[X]-\underline{m})(\overline{m}+\underline{m}) = (\overline{m}-\mu)(\mu-\underline{m}) + \mu^2.$$

Thus, equation (B.2) holds. As a result, equation (B.1) becomes

$$\min_{\substack{\underline{m}\leq X\leq \overline{m}\\ \mathbb{E}[X]\geq w_1}} \frac{\left(\mathbb{E}[X]-w_2\right)^2}{\mathsf{Var}[X]} = \min_{w_1\leq \mu\leq \overline{m}} g(\mu), \tag{B.5}$$

where

$$g(\mu) = \frac{(\mu - w_2)^2}{(\overline{m} - \mu)(\mu - \underline{m})}.$$

Note that the function $g(\cdot)$ is increasing for any μ satisfying $\underline{m} \leq w_2 \leq \mu \leq \overline{m}$ as

$$\frac{dg(\mu)}{d\mu} = \frac{(\mu - w_2)\left[(\overline{m} - \mu)(\mu + w_2 - 2\underline{m}) + (\mu - \underline{m})(\mu - w_2)\right]}{(\overline{m} - \mu)^2(\mu - \underline{m})^2} \geq 0.$$

Since $w_2 \leq w_1$, this would then imply that the minimum on the right hand side of (B.5) is obtained at $\mu = w_1$. This will complete the proof.

References

Adler, I., 2013. The equivalence of linear programs and zero-sum games. International Journal of Game Theory 42 (1), 165–177.

Broniatowski, M., Leorato, S., 2006. An estimation method for the neyman chi-square divergence with application to test of hypotheses. Journal of multivariate analysis 97 (6), 1409–1436.

Budinich, M., Fortnow, L., 2011. Repeated matching pennies with limited randomness. In: Proceedings of the 12th ACM conference on Electronic commerce. ACM, ACM, pp. 111–118.

Cicalese, F., Gargano, L., Vaccaro, U., 2017. How to find a joint probability distribution of minimum entropy (almost), given the marginals. arXiv preprint arXiv:1701.05243.

Cover, T. M., Thomas, J. A., 2012. Elements of information theory. John Wiley and Sons.

Dantzig, G. B., 1951. A proof of the equivalence of the programming problem and the game problem. Activity analysis of production and allocation (13), 330–338.

Fang, S.-C., Rajasekera, J. R., Tsao, H.-S. J., 2012. Entropy optimization and mathematical programming. Vol. 8. Springer Science and Business Media.

Gossner, O., Vieille, N., 2002. How to play with a biased coin? Games and Economic Behavior 41 (2), 206–226.

Han, T. S., 2003. Information-spectrum methods in information theory. Vol. 50. Springer -Verlag Berlin Heidelberg.

Hubek, P., Naor, M., Ullman, J., 2016. When can limited randomness be used in repeated games? Theory of Computing Systems 59 (4), 722–746.

Kapur, J. N., Baciu, G., Kesavan, H. K., 1995. The minmax information measure. International Journal of Systems Science 26 (1), 1–12.

Kocaoglu, M., Dimakis, A. G., Vishwanath, S., Hassibi, B., 2016. Entropic causal inference. arXiv preprint arXiv:1611.04035.

Kovacevic, M., Stanojevic, I., Senk, V., 2012. On the hardness of entropy minimization and related problems. In: Information Theory Workshop (ITW), 2012 IEEE. IEEE, IEEE, pp. 512–516.

Kovaevi, M., Stanojevi, I., enk, V., 2015. On the entropy of couplings. Information and Computation 242, 369–382.

Narahari, Y., 2014. Game Theory and Mechanism Design. Vol. 4. World Scientific.

Nash, J. F., 1950. Equilibrium points in n-person games. Proceedings of the national academy of sciences 36 (1), 48–49.

Neyman, A., Okada, D., 2000. Repeated games with bounded entropy. Games and Economic Behavior 30 (2), 228–247.

Neyman, A., Okada, D., 2009. Growth of strategy sets, entropy, and nonstationary bounded recall. Games and Economic Behavior 66 (1), 404–425.

Shor, P. W., 2004. Equivalence of additivity questions in quantum information theory. Communications in Mathematical Physics 246 (3), 453–472.

Slepian, D., Wolf, J., 1973. Noiseless coding of correlated information sources. IEEE Transactions on information Theory 19 (4), 471–480.

Srikanth, M., Kesavan, H., Roe, P., 2003. Computation of the MinMax measure. Springer, pp. 239–251.

Watanabe, S., 1981. Pattern recognition as a quest for minimum entropy. Pattern Recognition 13 (5), 381–387.

Yassaee, M. H., Aref, M. R., Gohari, A., 2014. Achievability proof via output statistics of random binning. IEEE Transactions on Information Theory 60 (11), 6760–6786.